

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MASSACHUSETTS**

MASSACHUSETTS INSTITUTE OF)
TECHNOLOGY)
)
)
Plaintiff,)
)
)
v.) Civil Action No. 05-10990-DPW
)
HARMAN INTERNATIONAL INDUSTRIES,)
INCORPORATED,)
)
)
Defendant.)
)
)

**HARMAN'S MEMORANDUM IN SUPPORT OF ITS
MOTION FOR SUMMARY JUDGMENT OF
UNENFORCEABILITY DUE TO MIT'S INEQUITABLE CONDUCT**

TABLE OF CONTENTS

	Page
I. FACTUAL BACKGROUND.....	2
A. MIT's Conduct That Led To The Patent-In-Suit.....	2
B. MIT Compounded Its Deceit By Withholding Critical Information In This Action.....	9
II. ARGUMENT	10
A. MIT Committed Inequitable Conduct When It Affirmatively Misrepresented And Concealed The Availability Of The Davis Thesis.	12
1. The Thesis Information That MIT Misrepresented and Withheld Was Material.....	12
2. MIT Intended to Deceive the PTO by Misrepresenting and Withholding the Thesis Information.....	13
B. MIT Committed Inequitable Conduct When It Withheld Information About The 50 People Who Used Completed Back Seat Driver Systems Before The Critical Date.....	15
1. The Information MIT Withheld Information About the Public Uses of The Back Seat Driver Was Material.	15
2. MIT's Conduct Regarding the Public Uses Was Intentionally Deceptive.....	18
III. CONCLUSION.....	19

TABLE OF AUTHORITIES

Federal Cases

<i>Agfa Corp. v. Creo Prods. Inc.</i> , 451 F.3d 1366 (Fed. Cir. 2006)	10, 11
<i>Amgen, Inc. v. Hoechst Marion Roussel, Inc.</i> , 314 F.3d 1313 (Fed. Cir. 2003)	14
<i>Anderson v. Liberty Lobby, Inc.</i> , 477 U.S. 242 (1986)	11
<i>Avia Group Int'l, Inc. v. L.A. Gear Cal., Inc.</i> , 853 F.2d 1557 (Fed. Cir. 1988)	11
<i>Barmag Barmer Maschinenfabrik AG v. Murata Mach., Ltd.</i> , 731 F.2d 831 (Fed. Cir. 1984)	12
<i>Baxter Int'l, Inc. v. Cobe Labs., Inc.</i> , 88 F.3d 1054 (Fed. Cir. 1996)	17
<i>Buehler AG v. Ocrim, S.p.A.</i> , 836 F. Supp. 1291 (N.D. Tex. 1992)	10
<i>Cargill Inc. v. Canabra Foods, Ltd.</i> , 476 F.3d 1359 (Fed. Cir. 2007)	2, 8, 14, 17, 19
<i>Celotex Corp. v. Catrett</i> , 477 U.S. 317 (1986)	11
<i>Conroy v. Reebok Int'l, Ltd.</i> , 14 F.3d 1570 (Fed. Cir. 1994)	11
<i>Constant v. Advanced Micro-Devices, Inc.</i> , 848 F. 2d 1560 (Fed. Cir. 1988)	5, 15
<i>Critikon, Inc. v. Becton Dickinson Vascular Access, Inc.</i> , 120 F.3d 1253 (Fed. Cir. 1997)	14, 18
<i>Digital Control Inc. v. The Charles Mach. Works</i> , 437 F.3d 1309 (Fed. Cir. 2006)	1, 17
<i>Dippin' Dots, Inc. v. Mosey et al.</i> , 476 F.3d 1337 (Fed. Cir. 2007)	1, 2, 12, 13, 17, 18
<i>Egbert v. Lippmann</i> , 104 U.S. 333 (1881)	1, 5, 16, 17

TABLE OF AUTHORITIES (cont'd)

	Page(s)
<i>Espeed, Inc. et al. v. Brokertec USA, L.L.C.</i> , No. 2006-1385, 2007 WL 817665, (Fed. Cir. March 20, 2007).....	7, 13
<i>Ferring B.V. v. Barr Labs., Inc.</i> , 437 F.3d 1181 (Fed. Cir. 2006)	10, 11, 14, 18
<i>Gardco Mfg., Inc. v. Herst Lighting Co.</i> , 820 F.2d 1209 (Fed. Cir. 1987)	10
<i>Hall v. Macneale</i> , 107 U.S. 90 (1883)	16
<i>Harrington Mfg. Co. v. Powell Mfg. Co.</i> , 815 F.2d 1478 (Fed. Cir. 1986)	1, 5, 12, 17
<i>In re '639 Patent Litigation</i> , 154 F. Supp. 2d 157 (D. Mass. 2001).....	13
<i>In re Epstein</i> , 32 F.3d 1559 (Fed. Cir. 1994)	16
<i>LaBounty Mfg., Inc. v. U.S. Int'l Trade Comm'n</i> , 958 F.2d 1066 (Fed.Cir.1992)	8, 19
<i>Li Second Family Ltd. P'ship v. Toshiba Corp.</i> , 231 F.3d 1373 (Fed. Cir. 2000)	13, 17
<i>Lough v. Brunswick Corp.</i> , 86 F.3d 1113 (Fed. Cir. 1996)	16
<i>Mass. Inst. of Tech. v. AB Fortia</i> , 774 F.2d 1104 (Fed. Cir. 1985)	8, 13, 15
<i>Matsushita Elec. Indus. Co. v. Zenith Radio Corp.</i> , 475 U.S. 574 (1986)	11
<i>Merck & Co., Inc. v. Danbury Pharmacal, Inc.</i> , 873 F.2d 1418 (Fed. Cir. 1989)	18
<i>Minn. Mining and Mfg. Co. v. Appleton Papers Inc.</i> , 35 F. Supp. 2d 1138 (D. Minn. 1999).....	17
<i>Molins PLC v. Textron, Inc.</i> , 48 F.3d 1172 (Fed. Cir. 1995)	2, 12, 18
<i>Monsanto Co. v. Bayer Bioscience N.V.</i> , 363 F.3d 1235 (Fed. Cir. 2004)	2

TABLE OF AUTHORITIES (cont'd)

	Page(s)
<i>New Railhead Mfg., L.L.C. v. Vermeer Mfg. Co.,</i> 298 F.3d 1290 (Fed. Cir. 2002)	16
<i>Paragon Podiatry Lab., Inc. v. KLM Labs., Inc.,</i> 984 F.2d 1182 (Fed. Cir. 1993)	10, 11, 15, 18
<i>Precision Instrument Mfg. Co. v. Auto. Maint. Mach. Co.,</i> 324 U.S. 806 (1945)	2
<i>Refac Intern., Ltd. v. Lotus Dev. Corp.,</i> 81 F.3d 1576 (Fed. Cir. 1996)	15, 18
<i>Semiconductor Energy Lab. Co. v. Samsung Elecs. Co.,</i> 204 F.3d 1368 (Fed. Cir. 2000)	7
<i>Zubulake v. UBS Warburg LLC,</i> \br/>220 F.R.D. 212 (2d Cir. 2003).....	9
State Cases	
35 U.S.C. §102	5, 14, 17
Cases	
57 FED. REG. 2021 (Jan. 17, 1992)	17
Federal Statutes	
37 C.F.R. § 1.56 (1991)	16
37 C.F.R. § 1.56 (2006)	12, 13, 17
FED. R. CIV. P. 56(c)	11
FED. R. CIV. P. 56(e)	11

Plaintiff Massachusetts Institute of Technology (“MIT”) faced tremendous pressure during prosecution of the sole patent-in-suit, U.S. Patent No. 5,177,685 (the “’685 Patent”). A major donor to the MIT Media Lab who also financed the Back Seat Driver project that led to the ’685 Patent, NEC Home Electronics (“NEC”), was disgruntled [[REDACTED]]. Ex. 6 at 7388. [[REDACTED]]. Ex. 6. MIT knew that having [[REDACTED]] jeopardized the significant flow of sponsorship dollars on which MIT’s Media Lab relied for its existence. Ex. 8 at 7392.

MIT needed to placate its upset corporate sponsor, and the only sure way to do so was to make sure the ’685 Patent issued. So, MIT misrepresented key facts and failed to disclose others to the United States Patent and Trademark Office (“PTO”) during prosecution of the ’685 Patent. Specifically, MIT affirmatively misrepresented and concealed the availability of the inventor’s thesis after the PTO Examiner expressly found it was prior art that invalidated each and every claim of the ’685 Patent. *See Section II.A. below; see also Dippin’ Dots, Inc. v. Mosey et al.*, 476 F.3d 1337, 1337 (Fed. Cir. 2007) (“A patent may be rendered unenforceable for inequitable conduct if an applicant, with intent to mislead or deceive the examiner, fails to disclose material information or submits materially false information to the PTO during prosecution.” (*citing Digital Control Inc. v. The Charles Mach. Works*, 437 F.3d 1309, 1313 (Fed. Cir. 2006))). Separately, MIT withheld information and deceived the PTO about 50 public uses of completed systems that were already known to be working satisfactorily—50 public uses when just one public use bars the patent under Section 102(b) of the Patent Act. *See Section II.B. below; see also Harrington Mfg. Co. v. Powell Mfg. Co.*, 815 F.2d 1478, 1481 (Fed. Cir. 1986) (invalidating public use where one of three inventors demonstrated a prototype of the invention to one person); *Egbert v. Lippmann*, 104 U.S. 333, 336 (1881). In both cases, MIT intended to deceive the PTO in order to secure allowance of a patent to which it is not entitled.

“[P]atent monopolies [must] spring from backgrounds free from fraud or other inequitable conduct.” *Precision Instrument Mfg. Co. v. Auto. Maint. Mach. Co.*, 324 U.S. 806, 816 (1945). Not only are applicants “required to prosecute patent applications in the PTO with candor, good faith, and honesty”—any intentional “breach of this duty constitutes inequitable conduct”—applicants have an “uncompromising duty” to inform the PTO of all facts concerning possible inequitable acts. *Molins PLC v. Textron, Inc.*, 48 F.3d 1172, 1178 (Fed. Cir. 1995); *Precision Instrument*, 324 U.S. at 818.

Here, the clear and convincing evidence—indeed undisputed—evidence shows that MIT intentionally breached its duties of candor and disclosure not once, but twice.¹ Just one breach is sufficient with respect to materiality and intent to establish inequitable conduct, thereby warranting a holding that the ’685 Patent is unenforceable. *See Cargill Inc. v. Canabra Foods, Ltd.*, 476 F.3d 1359, 1364 (Fed. Cir. 2007) (Once thresholds of both materiality and intent are met, the Court then “balance[s] the equities to determine whether the patentee has committed inequitable conduct that warrants holding the patent unenforceable.” (quoting *Monsanto Co. v. Bayer Bioscience N.V.*, 363 F.3d 1235, 1239 (Fed. Cir. 2004))); *Dippin’ Dots*, 476 F.3d at 1345-46. Taken together, the combination of acts compels such a finding and merits dismissal of this action before the parties and this Court waste any additional resources.

I. FACTUAL BACKGROUND

A. MIT’s Conduct That Led To The Patent-In-Suit.

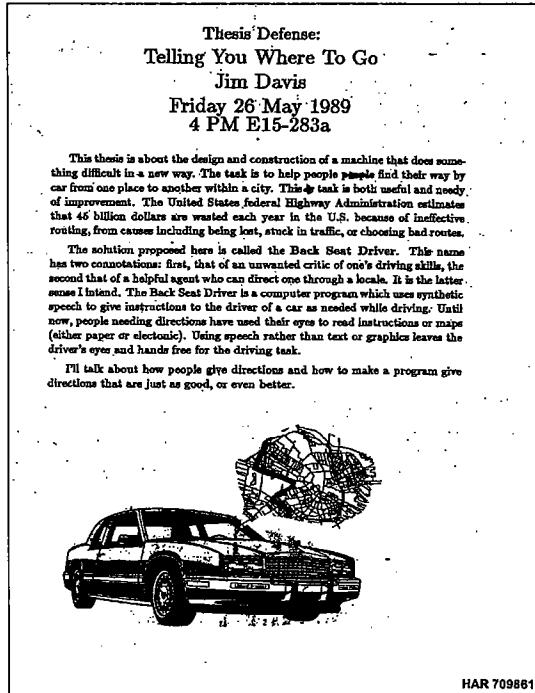
In the late 1980’s, Jim Davis was a graduate student at MIT working within the Speech Research Group of the MIT Media Lab under his faculty advisor, Chris Schmandt, who also

¹ Harman has identified other instances of inequitable conduct, which are not presented in this motion, in its interrogatory responses and expert reports. Harman expressly reserves the right to raise these additional instances of inequitable conduct where appropriate.

served as director of that Group. Harman's Local Rule 56.1 Statement of Undisputed Facts ("SOF") 1. By early 1988, Davis and Schmandt were working on a project called the Back Seat Driver that involved automobile navigation using spoken directions. SOF 2. The Back Seat Driver became Davis' thesis for his doctorate degree and eventually led to the patent application that issued as the '685 Patent. SOF 3. A subsidiary of the large, Japanese corporation NEC funded the project. SOF 4.

By [[REDACTED]] and had [[REDACTED]] on the public streets around the Boston area. Ex. 18 at 108; SOF 39. In fact, Davis and Schmandt were [[REDACTED]] and had already done so. Ex. 18 at 108, 112; SOF 40. Between [[REDACTED]], at least [[REDACTED]] used the Back Seat Driver to navigate the public streets of Boston. Ex. 19 at 173; SOF 41. Among the subjects who used the Back Seat Driver were [[REDACTED]], and several General Motors' employees, none of whom signed any type of confidentiality agreement or were bound by any confidentiality obligation. SOF 42-43. Indeed, there is no evidence that any of the 50 people who were shown the Back Seat Driver ever pledged any confidentiality at all. SOF 41, 43. Yet, they openly used the Back Seat Driver all over Boston, while MIT stored the system in a public garage. SOF 44.

By May 26, 1989, Davis was prepared to defend his thesis and invited the public to attend:



Ex. 5; SOF 12.

Like the Back Seat Driver itself, Davis' thesis was no secret. [[REDACTED]] while Davis was working on his thesis, and even after he finished drafting it, Davis [[REDACTED]] SOF 5. That is exactly what MIT did. On [[REDACTED]], Davis and Schmandt sent a copy of the completed thesis ("certified by Nicholas P. Negroponte," the then Director of the MIT Media Lab) to a Mr. Rittmueller, an NEC employee at the time. SOF 6. It was not marked confidential in any way. SOF 7. Davis himself sent another copy of his thesis (unsigned but bearing an August 4, 1989 date for Davis' signature) to a Bell Labs employee, Lynn Streeter, without any confidentiality designation and without any restrictions on her use of it. See SOF 8. [[REDACTED]] SOF 9. The words, lines, and pages of the copies sent to Dr. Streeter and Mr. Rittmueller are identical to the copy that MIT later submitted to the PTO with its patent application; only the signature differs on the title page provided to Mr. Rittmueller:

Back Seat Driver: voice assisted automobile navigation
by
James Raymond Davis
B.S.A.D., Massachusetts Institute of Technology (1977)
Submitted to the Media Arts and Sciences Section
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
September 1989
©Massachusetts Institute of Technology 1989
All Rights Reserved

Signature of Author Media Arts and Sciences Section
August 4, 1989
Certified by Nicholas P. Negroponte
Professor of Media Technology
Thesis Supervisor
Accepted by Stephen A. Belotti
Chairman, Departmental Committee on Graduate Students

STREETER 880

Back Seat Driver: voice assisted automobile navigation
by
James Raymond Davis
B.S.A.D., Massachusetts Institute of Technology (1977)
Submitted to the Media Arts and Sciences Section
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
September 1989
©Massachusetts Institute of Technology 1989
All Rights Reserved

Signature of Author Media Arts and Sciences Section
August 4, 1989
Certified by Nicholas P. Negroponte
Professor of Media Technology
Thesis Supervisor
Accepted by Stephen A. Belotti
Chairman, Departmental Committee on Graduate Students

RITTMUELLER 880

Back Seat Driver: voice assisted automobile navigation
by
James Raymond Davis

B.S.A.D., Massachusetts Institute of Technology (1977)
Submitted to the Media Arts and Sciences Section
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
September 1989
©Massachusetts Institute of Technology 1989
All Rights Reserved

Signature of Author Media Arts and Sciences Section
August 4, 1989
Certified by Nicholas P. Negroponte
Professor of Media Technology
Thesis Supervisor
Accepted by Stephen A. Belotti
Chairman, Departmental Committee on Graduate Students

RITTMUELLER 880

See Ex. 4 (Streeter); Ex. 3 (Rittmueller); Ex. 2 at 111 (PTO); SOF 11.

Davis' thesis was available to other interested members of the public. [[REDACTED]]²

SOF 10. Davis' openness regarding his research project was consistent with—indeed, prescribed by—[[REDACTED]]. SOF 13.

Davis signed and submitted his thesis on August 4, 1989. SOF 14. One year and five days later, on August 9, 1990, MIT filed a patent application, based on Davis' thesis, that became the '685 Patent. SOF 15. Section 102(b) of the Patent Act precludes the PTO from granting a patent if the subject matter sought to be claimed was “described in a printed publication...or in public use...more than one year prior to the date of the application for patent in the United States.” 35 U.S.C. §102. Thus any publication or public use before August 9, 1989 (the “critical date”) bars the '685 Patent. See *Constant v. Advanced Micro-Devices, Inc.*, 848 F. 2d 1560, 1569 (Fed. Cir. 1988) (for publication “there is no requirement to show that particular members of the public actually received the information” so long as “interested members of the relevant public could obtain the information if they wanted to”); *Harrington*, 815 F.2d at 1481; *Egbert*, 104 U.S. at 336.

² [[REDACTED]]

After the Back Seat Driver application had been pending about a year without any substantive examination on the merits, [[REDACTED]]:

[[REDACTED]]

Ex. 6 at 7388; SOF 17. [[REDACTED]]:

[[REDACTED]]

Ex. 6 at 7387; SOF 18. [[REDACTED]]. Ex. 6 at 7388; SOF 19. Acutely aware that its research sponsors, particularly NEC, saw intellectual property rights as an important benefit of their agreement to commit hundreds of thousands of dollars of research grants to MIT's Media Lab, [[REDACTED]]:

[[REDACTED]]

Ex. 8 at 7392; SOF 20.

Shortly thereafter, on November 8, 1991, the PTO took some action and rejected each and every pending claim because the Davis thesis was publicly available more than one year before MIT filed the patent application:

~~Claims 1-58 are rejected under 35 U.S.C. § 102(e)
as being clearly anticipated by the Ph.D. thesis of
James Raymond Davis.~~

November 4, 1991 Office Action, Ex. 2 at 442; SOF 21.

[[REDACTED]]:

[[REDACTED]]

Ex. 7 at 7389; SOF 22. To the PTO, MIT expressly acknowledged that it clearly understood the Examiner's reasons for rejection:

The examiner has rejected claims 1-58 under 35 U.S.C. 102(e) as being anticipated by the Ph.D. thesis of J.R. Davis. During a telephone conversation with the examiner, the examiner stated that the reason for the rejection was that the title page of the thesis bears a submission date of August 4, 1989, more than one year before the filing date of the present application.

MIT's May 5, 1992 Response, Ex. 2 at 803; SOF 23 (emphasis added).

In responding to the Examiner, MIT did not challenge the materiality of Davis' thesis or the Examiner's determination that the thesis anticipated each claim. Rather, MIT submitted a new title page for the thesis that was stamped by MIT's library and, based on that library stamp, challenged the date the thesis became available to the public:

August 4 is the date that the thesis was signed, and not the date on which the thesis became available to the public. M.I.T. does not generally catalog and shelve theses until several months after the official date of submission. Enclosed is a copy of the title page of the M.I.T. library's copy of the thesis, which bears a date of February 27, 1990. Therefore, the thesis did not become available to the public more than a year before the filing date of the present application, and is therefore not 102 art with respect to the present application.

MIT's May 5, 1992 Response, Ex. 2 at 803; SOF 24-27 (emphasis added). While MIT affirmatively represented that "the thesis did not become available to the public more than a year before the filing date of the present application," MIT knew and actively concealed: (i) that Davis could disseminate his thesis on request; (ii) that Davis actually distributed his thesis and that he did so more than once; and (iii) that Davis publicly defended his thesis. SOF 28-30. Not only was MIT's representation false in light of these facts, it misled the Examiner into believing that the only access to Davis' thesis was through the MIT library, which Davis, Schmandt and, hence, MIT all knew to be wrong. *Espeed, Inc. et al. v. Brokertec USA, L.L.C.*, No. 2006-1385, 2007 WL 817665, at *7 (Fed. Cir. March 20, 2007) (attached as Ex. 29) (finding a misrepresentation to the PTO material where it "left the examiner with the impression that the examiner did not need to conduct any further...investigation" (quoting *Semiconductor Energy Lab. Co. v. Samsung Elecs. Co.*, 204 F.3d 1368, 1377 (Fed. Cir. 2000))). Even if MIT believed

these were not public disclosures, MIT was not entitled to act as judge and jury by making that determination itself; rather, MIT was required to disclose the information and let the PTO decide the issue. *See Cargill*, 476 F.3d at 1367 (quoting *LaBounty Mfg., Inc. v. U.S. Int'l Trade Comm'n*, 958 F.2d 1066, 1076 (Fed. Cir. 1992) (any doubts as to materiality must be “resolved in favor of disclosure”)).

MIT kept all of this information from the PTO even though, or perhaps because, it had recently been the plaintiff in a case where the Federal Circuit held that a reference was § 102 art because the inventor “orally presented” it and distributed copies of it “on request, without any restrictions...more than one year before the filing date of the...patents.” *Mass. Inst. of Tech. v. AB Fortia*, 774 F.2d 1104, 1108–09 (Fed. Cir. 1985). In *Fortia*, the court went on to invalidate the patents at issue. *Id.* at 1109. When confronted with the same situation and facing the same outcome here, MIT chose to avoid the problem altogether by making false statements and misleading the Examiner.

Separately, MIT never told the Examiner about the [[REDACTED]] who used the Back Seat Driver to navigate the public streets of Boston more than a year before MIT filed the patent application. SOF 45-46. MIT never told the PTO that those public uses embodied the subject matters that were being claimed, or that the uses continued even after the claimed subject matter was reduced to practice. SOF 47-49. MIT withheld the timing of those public uses, their extent, the details surrounding them, and, perhaps most importantly, the lack of confidentiality regarding them. SOF 50-53.

Even when subsequently disclosing information that MIT did “want to be considered and made of record,” MIT withheld material information about the thesis availability and distribution as well as material information about 50 public uses of the Back Seat Driver. SOF 46. With

more than 800 U.S. patents already issued in its name, MIT was a seasoned patent applicant when it filed and prosecuted the application that led to the '685 Patent. *See, e.g.*, SOF 16. MIT no doubt understood the PTO rules and the dire consequences of disclosure. MIT's lack of candor resulted in a Notice of Allowance of the '685 Patent on June 30, 1992. June 30, 1992 Notice of Allowance, Ex. 2 at 808; SOF 32. [[REDACTED]]. Ex. 9 at 7404; SOF 33.

B. MIT Compounded Its Deceit By Withholding Critical Information In This Action.

MIT continued its deception in this litigation by concealing, in some cases actively, the documents evidencing the information MIT misrepresented or withheld from the PTO. Importantly, all of those documents bear a 1989 date, and MIT contends it had an "eye toward litigation" in 1989 giving rise to privilege and thus triggering its obligation to preserve documents. *See* Ex. 27; SOF 34-38; *see also Zubulake v. UBS Warburg LLC*, 220 F.R.D. 212, 216 (2d Cir. 2003) ("The obligation to preserve evidence arises when the party has notice that the evidence is relevant to litigation or when a party should have known that the evidence may be relevant to future litigation." (citations omitted)).

MIT did not produce the document showing Davis' Thesis Defense. HAR 709861, Ex. 5; SOF 37. Harman did. Harman discovered that document serendipitously through a corporate partner in its counsel's law firm who happened to work on the Back Seat Driver while an undergraduate student at MIT. MIT has not produced a single document evidencing Davis' defense of his thesis.

Nor did MIT produce the documents it authored and sent to Mr. Rittmueller in 1989. *See, e.g.*, Ex. 3; Ex. 18; Ex. 19; SOF 36. On the contrary, MIT improperly asserted privilege and attempted to conceal their mere existence from Harman in this litigation. *See* Ex. 28; SOF 34. Even after this Court ordered MIT to produce certain documents, MIT did not produce (i) the

thesis MIT sent to Mr. Rittmueller, or (ii) [[REDACTED]], or (iii) [[REDACTED]]. Ex. 3; Ex. 18; Ex. 19. Harman obtained these documents when it subpoenaed Mr. Rittmueller, whom MIT's litigation counsel here represented. *See* SOF 36.

MIT also did not produce the thesis it sent to Dr. Streeter in 1989. Ex. 4. Harman obtained this evidence only after it subpoenaed Dr. Streeter. *See* SOF 35. That MIT has been hiding the facts regarding the availability of Davis' thesis and the 50 public uses for more than 15 years speaks volumes about MIT's deceptive intent.

II. ARGUMENT

An affirmative defense of unenforceability due to inequitable conduct is amenable to summary judgment in the absence of a genuine issue of material fact. *Agfa Corp. v. Creo Prods. Inc.*, 451 F.3d 1366, 1380 (Fed. Cir. 2006) (affirming D. Mass. Court's grant of summary judgment that withholding a reference directly adverse to patentee's arguments for patentability constituted inequitable conduct); *Paragon Podiatry Lab., Inc. v. KLM Labs., Inc.*, 984 F.2d 1182, 1192–93 (Fed. Cir. 1993) (granting summary judgment of inequitable conduct where patentee failed to disclose pre-critical date sales and submitted false and misleading affidavits to the PTO); *Ferring B.V. v. Barr Labs., Inc.*, 437 F.3d 1181, 1195 (Fed. Cir. 2006) (affirming summary judgment of unenforceability); *accord Buehler AG v. Ocrim, S.p.A.*, 836 F. Supp. 1291, 1299 (N.D. Tex. 1992), *aff'd*, 34 F.3d 1080 (Fed. Cir. 1994).

The inquiry into whether MIT engaged in inequitable conduct is a *purely equitable one* for which there is no right to a jury. *See Agfa*, 451 F.3d at 1375 (re-affirming the holding in *Gardco Mfg., Inc. v. Herst Lighting Co.*, 820 F.2d 1209, 1212 (Fed. Cir. 1987) (finding no right to a jury trial because “inequitable conduct” is...purely equitable in nature)). Claim construction has no bearing on the facts presented here, making this issue particularly ripe for summary disposition in advance of claim construction and other dispositive motions. Summary

judgment on this purely equitable question would dispense with this case altogether.³ See *Ferring*, 437 F.3d at 1195 (granting summary judgment of inequitable conduct and dismissing the case).

Summary judgment in Harman's favor is warranted because "there is no genuine issue as to any material fact and [Harman] is entitled to judgment as a matter of law." FED. R. CIV. P. 56(c); see also *Celotex Corp. v. Catrett*, 477 U.S. 317, 325 (1986) ("the burden on the moving party may be discharged by 'showing'...that there is an absence of evidence to support the nonmoving party's case"); *Conroy v. Reebok Int'l, Ltd.*, 14 F.3d 1570, 1575 (Fed. Cir. 1994) ("The grant of summary judgment is appropriate in a patent case where the standards set forth in Rule 56(c) are satisfied." (citing *Paragon*, 984 F.2d at 1190)). To avoid summary judgment, MIT must do more than simply challenge a fact, it must come forward with "specific facts showing that there is a genuine issue for trial." FED. R. CIV. P. 56(e); see also *Avia Group Int'l, Inc. v. L.A. Gear Cal., Inc.*, 853 F.2d 1557, 1560 (Fed. Cir. 1988) (non-movant "must do more than merely raise some doubt as to the existence of a fact; evidence must be forthcoming from the nonmovant which would be sufficient to require submission to the jury of the dispute over the fact." (citing *Anderson v. Liberty Lobby, Inc.*, 477 U.S. 242, 250 (1986))); *Matsushita Elec. Indus. Co. v. Zenith Radio Corp.*, 475 U.S. 574, 586–87 (1986) (non-movant "must do more than simply show that there is some metaphysical doubt as to the material facts...the nonmoving party must come forward 'with specific facts showing that there is a genuine issue for trial.'") (quoting FED. R. CIV. P. 56(e)); *Barmag Barmer Maschinenfabrik AG v. Murata Mach., Ltd.*, 731 F.2d

³ Should the Court find this issue cannot be disposed on this motion, Harman respectfully requests that the Court sever inequitable conduct from the rest of the case and conduct a short bench trial on that issue before proceeding with a jury trial on other issues to avoid further unnecessary litigation costs. See *Agfa*, 451 F.3d at 1371 ("[T]he conduct-of-the-applicant-in-the-PTO issue raised in the nonjury trial and the separated infringement/validity issues are distinct and without commonality either as claims or in relation to the underlying fact issues."); *Gardco*, 820 F.2d at 1211-13.

831, 836 (Fed. Cir. 1984) ("[T]he court may not simply accept a party's statement that a fact is challenged.")

There is no genuine issue of material fact that MIT intentionally and repeatedly violated its duty of candor during prosecution of the patent at issue in this case. MIT did so when it submitted false information to the PTO regarding the availability of Davis' thesis. *See Dippin' Dots*, 476 F.3d at 1345. MIT did so again when it failed to disclose information about the [[REDACTED]] who drove the Back Seat Driver around Boston more than one year before the patent application was filed. *See Harrington*, 815 F.2d at 1481; *see also* 37 C.F.R. § 1.56 (2006) ("The public interest is best served [when]...the [Patent] Office is aware of and evaluates the teachings of *all information* material to patentability." (emphasis added)). Either one demonstrates that MIT did not conduct itself with the "candor, good faith and honesty" essential to the patent system. *Molins*, 48 F.3d at 1178.

A. MIT Committed Inequitable Conduct When It Affirmatively Misrepresented And Concealed The Availability Of The Davis Thesis.

1. The Thesis Information That MIT Misrepresented and Withheld Was Material.

There is no question that the availability of the Davis thesis, which MIT misrepresented and about which MIT misled the Examiner, was material. The Examiner cited the thesis and referenced its availability as the basis for rejecting all pending claims during prosecution. Ex. 2 at 442; SOF 21. Statements regarding an invalidating reference are highly material as a matter of

law. *See* 37 C.F.R. § 1.56 (2006); *accord Espeed*, Ex. 29 at *7. MIT did not challenge the materiality then, and it cannot credibly challenge it now.⁴ SOF 24.

Besides, MIT did not simply omit information when it responded to the Examiner. MIT represented, in support of its patentability argument, that “the thesis did not become available to the public more than a year before the filing date of the patent application, and is therefore not 102 art with respect to the present application.” Ex. 2 at 803; SOF 25-27. That statement is demonstrably false: MIT could and did publicly distribute the thesis more than a year before the filing date of the patent application. SOF 5-8, 28-29; *see also Mass. Inst. of Tech.*, 774 F.2d at 1108-09. As such, it is *per se* material. 37 C.F.R. § 1.56 (2006) (statements that refute or contradict “a position [that] the applicant takes in (i) Opposing an argument of unpatentability relied on by the Office, or (ii) Asserting an argument of patentability” are *per se* material.); *accord Espeed*, Ex. 29 at *7.

2. MIT Intended to Deceive the PTO by Misrepresenting and Withholding the Thesis Information.

Given the unassailable materiality, the only issue here is whether MIT intended to deceive the PTO when it submitted false and misleading information about the availability of the Davis thesis. The simple answer is yes: “A court may infer deceptive intent from submission of false statements to the PTO to overcome prior art.” *In re '639 Patent Litigation*, 154 F. Supp. 2d 157, 188 (D. Mass. 2001) (*citing Li Second Family Ltd. P'ship v. Toshiba Corp.*, 231 F.3d 1373, 1381 (Fed. Cir. 2000)). And where materiality is high, as here, the requisite level of intent is

⁴ Any argument by MIT that the thesis was not a “printed publication” under § 102(b) of the Patent Act misses the point and is irrelevant. This Court need not decide that invalidity issue to grant summary judgment in Harman’s favor. *See Li Second Family Ltd. P'ship v. Toshiba Corp.*, 231 F.3d 1373, 1380 (Fed. Cir. 2000) (finding materiality when “a reasonable examiner would have considered the information important, not whether the information would conclusively decide the issue of patentability”). For inequitable conduct, Harman only must show that MIT, “with intent to mislead or deceive the examiner, fail[ed] to disclose material information or submit[ted] materially false information to the PTO during prosecution.” *Dippin’ Dots*, 476 F.3d at 1345.

lower. See *Cargill*, 476 F.3d at 1364 (“[t]he more material the omission or the misrepresentation, the lower the level of intent required to establish inequitable conduct, and vice versa.” (quoting *Critikon, Inc. v. Becton Dickinson Vascular Access, Inc.*, 120 F.3d 1253, 1256 (Fed. Cir. 1997))); *Amgen, Inc. v. Hoechst Marion Roussel, Inc.*, 314 F.3d 1313, 1358 (Fed. Cir. 2003) (“a lesser quantum of evidence of intent is necessary when the omission or misrepresentation is highly material”). A patentee “facing a high level of materiality and clear proof that it knew or should have known of that materiality, can expect to find it difficult to establish ‘subjective good faith’ sufficient to prevent the drawing of an inference of intent to mislead.” *Critikon*, 120 F.3d at 1257; *Ferring*, 437 F.3d at 1191.

For MIT, patents and intellectual property were and are marketing tools used to lure Media Lab sponsors and provide such sponsors with at least the appearance of some tangible [[REDACTED]] in return for their generous sponsorships. Without research sponsorships, there is no Media Lab. The immediate benefit MIT reaps from filing patent applications and securing issued patents is the continuity of research dollars flowing from happy sponsors. Thus, MIT receives significant, tangible benefit from the mere issuance of patents.

Regarding the Back Seat Driver, MIT’s generous sponsor was already upset when the Examiner presented MIT with a statutory bar problem. SOF 18. MIT must have felt considerable pressure to see the Back Seat Driver issue as a patent in order to placate [[REDACTED]] that had [[REDACTED]] SOF 20. So much so, in fact, that MIT hid all of the NEC documents under baseless claims of privilege and forced Harman to move to compel before disclosing any documents [[REDACTED]]. SOF 34. Even then, MIT forced Harman to subpoena NEC for documents—documents authored by MIT—that MIT should have had in its possession. See *id.* MIT hid the documents evidencing MIT’s distribution of Davis’ thesis,

again forcing Harman to obtain from third parties what MIT should have had. *Id.*

If nothing else, “the inference [of an intent to mislead] arises [in part] from the inability of the examiner to investigate the facts.” *Refac Intern., Ltd. v. Lotus Dev. Corp.*, 81 F.3d 1576, 1582 (Fed. Cir. 1996) (*quoting Paragon*, 984 F.2d. at 1191). That Davis could and did distribute copies of his thesis demonstrates that the MIT library was not the only way for “interested members of the relevant public [to] obtain the information [Davis thesis] if they wanted to.” *Constant*, 848 F. 2d at 1569; SOF 5-8, 28-29. Yet, the Examiner was not made aware of any information whatsoever with which he could investigate the truth surrounding the Davis thesis, and there was no way for the PTO to ascertain that information independently. *See* SOF 27-30; *see also Mass. Inst. of Tech.*, 774 F.2d at 1108–09. The inventors, Davis and Schmandt, who assisted in prosecuting the patent, surely knew the truth, but robbed the Examiner of any chance to investigate the circulation of the thesis himself. *See* SOF 27-30. That is inequitable conduct.

B. MIT Committed Inequitable Conduct When It Withheld Information About The 50 People Who Used Completed Back Seat Driver Systems Before The Critical Date.

1. The Information MIT Withheld Information About the Public Uses of The Back Seat Driver Was Material.

MIT committed inequitable conduct again when it disclosed nothing about the [[REDACTED]] who drove the Back Seat Driver, very little information about any uses of the Back Seat Driver, and did so only in a very cursory fashion. In the patent specification, MIT noted only that “An actual working prototype of the Back Seat Driver has been implemented. It has successfully guided drivers unfamiliar with Cambridge, Mass. to their destinations.” Ex. 1 at 3:4-8; SOF 54. This statement reveals nothing about the timing, extent, lack of confidentiality, the exact subject matter involved, or the extent to which the uses embodied the same subject

matter attempting to be claimed. SOF 45-53. Quite the opposite—MIT downplayed these uses by adding that “[i]t is easy to foresee a practical implementation in the future,” (SOF 54), inferring that the uses did not themselves constitute a “practical implementation,” which MIT’s interrogatory responses belie. *See Ex. 10* (46 of the 58 claims of the ’685 Patent were reduced to practice before the critical date); SOF 47-48.⁵

There can be no genuine dispute that this withheld information about the public uses of the Back Seat Driver —the number of public uses, the identities of the third-party users, the lack of confidentiality, the claimed subject matters embodied in the public uses, that the invention was complete with respect to those claims, and that the subject matter being claimed was already reduced to practice at the time of the uses—was highly material.⁶ Information is material if “there is a substantial likelihood that a reasonable examiner would consider it important in deciding whether to allow the application to issue as a patent.” 37 C.F.R. § 1.56 (1991). Alternatively, “information is material to patentability when it is not cumulative...and...[i]t establishes, by itself or in combination with other information, a *prima facie* case of

⁵ Any claim by MIT that these uses were experimental is irrelevant and fails as a matter of law. They were not experimental because the subject matters recited in more than half of the ’685 Patent claims, including the sole independent claim, were reduced to practice “at least as early as June 1989.” Ex. 10; SOF 47-49; *New Railhead Mfg., L.L.C. v. Vermeer Mfg. Co.*, 298 F.3d 1290, 1299 (Fed. Cir. 2002) (“the patented method had been reduced to practice..., as a matter of law none of the subsequent uses of the method could be experimental.”). Even after that subject matter was reduced to practice, MIT and its inventor both admit that some of the [[REDACTED]] subjects who used the Back Seat Driver in public between [[REDACTED]], with no confidentiality obligation, continued their use, at least through [[REDACTED]]. SOF 41, 43, 47-49.

⁶ Any argument by MIT that these uses were not “public” such that MIT need not have disclosed any detailed information about them, is also misplaced. It is undisputed that the uses were on the public streets of Boston. SOF 44. The statutory phrase “public use” in Section 102(b) does not necessarily mean open and visible in the ordinary sense. *New Railhead*, 298 F. 3d at 1297 (citing *Lough v. Brunswick Corp.*, 86 F.3d 1113, 1119 (Fed. Cir. 1996); *Egbert*, 104 U.S. at 336. It has long been the rule that a use is still a “public use,” for purposes of Section 102(b), even though the use is such that the claimed invention is concealed from public view. *Sys. Mgmt. Arts*, 87 F. Supp. 2d at 270; *see also Egbert*, 104 U.S. at 336 (holding that a corset worn openly was in public use even though the invention was concealed within the corset); *Hall v. Macneale*, 107 U.S. 90 (1883) (holding that a safe mechanism was in public use even though the invention could not be seen without destroying the safe). The question is whether the invention was used in public, not whether the use disclosed the invention to the public. *In re Epstein*, 32 F.3d 1559, 1568 (Fed. Cir. 1994).

unpatentability of a claim....” 37 C.F.R. § 1.56 (2006).⁷

Even a single instance of public use before the critical date will bar patentability under Section 102(b). *See Harrington*, 815 F.2d at 1481; *Egbert*, 104 U.S. at 336. Here, the great number of uses compounds the materiality. *See Egbert*, 104 U.S. at 336 (“The use of a great number may tend to strengthen the proof, but one well-defined case of such use is just as effectual to annul the patent as many.”) Failing to disclose the public uses deprived the Examiner of considering the “totality of the circumstances in conjunction with the policies underlying the public use bar.” *Baxter Int'l, Inc. v. Cobe Labs., Inc.*, 88 F.3d 1054, 1058 (Fed. Cir. 1996). Although not critical here, that those public uses occurred under no limitation, restriction, or obligation of secrecy to the inventor indisputably affects patentability. *Baxter*, 88 F.3d at 1058-59 (affirming finding of public use where people, including co-workers of the inventor, who observed the invention in operation “were under no duty to maintain it as confidential”).

To grant summary judgment of inequitable conduct, the Court need not find that such uses were invalidating uses under 35 U.S.C. § 102. All that is necessary is that the Court find that the intentionally withheld information was material, *i.e.*, that a reasonable Examiner would have considered the material important. *See Li Second Family Ltd. P'ship*, 231 F.3d at 1380; *Dippin' Dots*, 476 F.3d at 1345-46. Even MIT’s patent prosecution attorney concedes that if [[REDACTED]] SOF 55.

⁷ The language of Rule 56 was revised in 1992 while MIT’s application was pending. 57 FED. REG. 2021 (Jan. 17, 1992) (revised Rule 56 applies to applications pending or filed after March 16, 1992). However, the Federal Circuit has made it clear that the revised Rule 56 “did not supplant the earlier ‘reasonable examiner’ standard.” *Cargill*, 476 F.3d at 1364. In other words, “if a misstatement or omission is material under the new Rule 56, it is material. Similarly, if a misstatement or omission is material under the ‘reasonable examiner’ standard...it is also material.” *Id.* (*citing Digital Control*, 437 F.3d at 1316). In any event, both versions of Rule 56 were in effect during the course of the prosecution of the ’685 Patent, and as such, both are applicable.

2. MIT's Conduct Regarding the Public Uses Was Intentionally Deceptive.

“[I]ntent may be inferred where a patent applicant knew, or should have known, that withheld information would be material to the PTO’s consideration of the patent application.” *Critikon*, 120 F.3d at 1256 (internal citations omitted). “Intent need not, and rarely can, be proven by direct evidence.” *Ferring*, 437 F.3d at 1191 (*quoting Merck & Co., Inc. v. Danbury Pharmacal, Inc.*, 873 F.2d 1418, 1422 (Fed. Cir. 1989)). Rather, it is “most often proven by a showing of acts, the natural consequences of which are presumably intended by the actor.” *Molins*, 48 F.3d at 1180; *Dippin’ Dots*, 476 F.3d at 1345 (Intent is generally “inferred from the facts and circumstances surrounding the applicant’s overall conduct.” (*citing Paragon*, 984 F.2d at 1189)). “Those who are not ‘up front’ with the PTO run the risk that, years later, a fact-finder might conclude that they intended to deceive.” *Molins*, 48 F.3d at 1182.

MIT’s prosecution attorney [[REDACTED]]. SOF 55. Yet, he made no effort to investigate the public uses of the “prototypes” disclosed in the specification. SOF 56. Both inventors, Davis and Schmandt, knew [[REDACTED]]. SOF 57-59. Indeed, both declared under penalty of perjury that they had “review[ed] and [understood] the contents of the...specification, including the claims” and each “acknowledge[d] the duty to disclose information which is material to the examination of [the] application in accordance with Title 37, Code of Federal Regulations, § 1.56.” Ex. 2 at 73-74. [[REDACTED]]. Ex. 14 at 74:24-75:3. Therefore, both Davis and Schmandt actually knew, or at the very least, should have known, of not only the pre-critical-date uses of the Back Seat Driver, but also the high materiality thereof under 37 C.F.R. § 1.56. See *Critikon*, 120 F.3d at 1256. Further evidencing deceptive intent, MIT prevented the Examiner from investigating the facts, choosing instead to do what it could to please its sponsor and get the Back Seat Driver patent issued. See *Refac Intern.*, 81 F.3d at 1582.

MIT could have erred on the side of disclosure; indeed, it should have. *See Cargill*, 476 F.3d at 1364. Instead, MIT “unilaterally withheld information that unquestionably would have been viewed as worthy of serious consideration by the PTO, and might have resulted in the patent[] not being issued.” *Id.* at 1365. MIT’s conduct was inequitable.

III. CONCLUSION

For these reasons, MIT engaged in at least two separate instances of inequitable conduct during prosecution of U.S. Patent No. 5,177,685, either of which renders the patent unenforceable. Harman respectfully requests that this Court enter summary judgment in Harman’s favor and dismiss this case in its entirety, with prejudice.⁸

⁸ Any argument by MIT that Harman’s motion is barred for failure to plead inequitable conduct with particularity is misplaced. MIT has been on notice of the factual bases supporting Harman’s inequitable conduct claim since at least June 25, 2004. Harman’s complaint filed in the Northern District of Illinois (and consolidated into this action in response to MIT’s motion to dismiss and transfer) stated the two then-known instances where MIT failed to disclose relevant information to the PTO during prosecution of the ’685 Patent. And when Harman uncovered documents (authored but not produced by MIT) supporting additional grounds for that claim, Harman provided full and complete notice of its new contentions to MIT. *See Ex. 27.* Indeed, MIT has litigated the inequitable conduct issue from the outset of this case by: (1) serving a contention interrogatory on this issue (which Harman answered and timely supplemented); (2) conducting and completing fact discovery on this issue; (3) proffering an expert who opined only on this issue and addressed *inter alia* the factual bases raised in the instant motion; (4) conducting and completing expert discovery on this issue.

Date: April 11, 2007

Respectfully submitted,

/s/ Courtney A. Clark

Robert J. Muldoon, Jr., BBO# 359480
James W. Matthews, BBO# 560560
Edward S. Cheng, BBO# 634063
Courtney A. Clark, BBO# 651381
SHERIN AND LODGEN, LLP
101 Federal Street
Boston, MA 02110

William A. Streff Jr., P.C.
Michelle A.H. Francis
Craig D. Leavell
Joanna Belle Gunderson
KIRKLAND & ELLIS LLP
200 East Randolph Drive
Chicago, Illinois 60601
(312) 861-2000 (phone)
(312) 861-2200 (fax)

*Attorneys for Defendant
Harman International Industries, Incorporated*

CERTIFICATE OF SERVICE

I hereby certify that this document filed through the ECF system will be sent electronically to the registered participants as identified on the Notice of Electronic Filing and paper copies will be sent to those indicated as non-registered participants on April 11, 2007.

/s/ Courtney A. Clark
Courtney A. Clark

EXHIBIT

1



US005177685A

United States Patent [19]

Davis et al.

[11] Patent Number: 5,177,685**[45] Date of Patent:** Jan. 5, 1993

[54] AUTOMOBILE NAVIGATION SYSTEM
USING REAL TIME SPOKEN DRIVING
INSTRUCTIONS

[75] Inventors: James R. Davis, North Cambridge;
Christopher M. Schmandt, Milton,
both of Mass.

[73] Assignee: Massachusetts Institute of
Technology, Cambridge, Mass.

[21] Appl. No.: 565,274

[22] Filed: Aug. 9, 1990

[51] Int. Cl.^s G01C 21/00

[52] U.S. Cl. 364/443; 340/988;
364/449; 364/453

[58] Field of Search 340/988, 989, 990, 995;
364/443, 444, 449, 450, 453, 436

[56] References Cited**U.S. PATENT DOCUMENTS**

4,139,889	2/1979	Ingels	340/989 X
4,697,281	9/1987	O'Sullivan	455/33 X
4,734,863	3/1988	Honey et al.	340/988 X
4,796,191	1/1989	Honey et al.	364/450
4,812,843	3/1989	Champion, III et al.	340/989 X
4,882,696	11/1989	Nimura et al.	
4,891,761	1/1990	Gray et al.	364/449 X
4,926,336	5/1990	Yamada	364/444

(List continued on next page.)

OTHER PUBLICATIONS

"Synthetic Speech for Real Time Direction-Giving", C. M. Schmandt et al., Digest of Technical Papers, International Conference on Consumer Electronics, Rosemont, Ill., Jun. 6-9, 1989.

"Synthetic Speech for Real Time Direction Giving", C. M. Schmandt et al., IEEE Transactions on Consumer Electronics, 35(3):649-653, Aug. 1989.

"The Back Seat Driver: Real Time Spoken Driving Instructions", J. R. Davis et al., Proceedings of the IEEE Vehicle Navigation and Information Systems Conference, Toronto, Canada, Sep. 1989.

"Back Seat Driver: Voice Assisted Automobile Navigation", by J. R. Davis, Ph.D. Thesis, Massachusetts Institute of Technology, Sep., 1989.

"CD-ROM Assisted Navigation Systems", by O. Ono

et al., Digest of Technical Papers, IEEE International Conference on Consumer Electronics, Rosemont, Ill., Jun. 8-10, 1988.

"Attention, Intentions, and the Structure of Discourse", by B. J. Grosz and C. L. Sidner (Computational Linguistics, 12(3):175-204, 1986).

"The Intonational Structuring of Discourse", by J. Hirschberg et al., Proceedings of the Association for Computational Linguistics, 136-144, Jul. 1986.

Automobile Electronic News, vol. 1, No. 16, "U.K. Picks GEC to Head Navigation Project", by James Fallon, Aug. 28, 1989.

"Softening of the Arteries", by Bruce Weber, The New York Times Magazine, p. 78, Aug. 26, 1990.

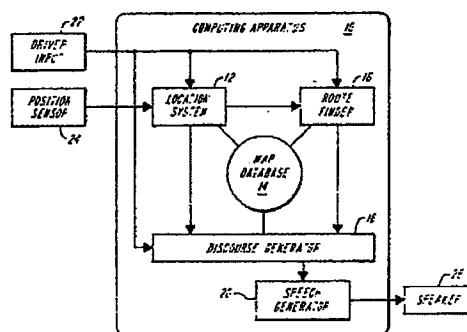
"Terminal Back Seat Driver", Technology Review, Jul., 1990, p. 10.

(List continued on next page.)

Primary Examiner—Parshotam S. Lal*Assistant Examiner*—Edward Pipala*Attorney, Agent, or Firm*—Choate, Hall & Stewart**[57] ABSTRACT**

An automobile navigation system which provides spoken instructions to the driver of an automobile to guide the driver along a route is disclosed. The heart of the system is a computing apparatus comprising a map database, route finding algorithms, a vehicle location system, discourse generating programs, and speech generating programs. Driver input means allows the driver to enter information such as a desired destination. The route finding algorithms in the computer apparatus calculate a route to the destination. The vehicle location system accepts input from a position sensor which measures automobile movement (magnitude and direction) continuously, and using this data in conjunction with the map database, determines the position of the automobile. Based on the current position of the automobile and the route, the discourse generating programs compose driving instructions and other messages according to a discourse model in real time as they are needed. The instructions and messages are sent to voice generating apparatus which conveys them to the driver.

58 Claims, 3 Drawing Sheets



5,177,685

Page 2

U.S. PATENT DOCUMENTS

4,937,751	6/1990	Nimura et al.	364/444 X
4,939,662	7/1990	Nimura et al.	340/990 X
4,951,211	8/1990	De Villerche	364/444
4,954,958	9/1990	Savage et al.	364/444
4,984,168	1/1991	Neukirchner et al.	364/449
4,992,947	2/1991	Nimura et al.	364/444
5,021,961	6/1991	Ross et al.	340/990 X
5,041,983	8/1991	Nakahara et al.	364/449
5,043,902	8/1991	Yokoyama et al.	364/449

OTHER PUBLICATIONS

"Taxi! Dynamic Cartographic Software for Training Cab Drivers", by M. Bosworth et al., Technical Report, Hunter College Department of Geology and Geography, (212)-772-4000, 1988 paper presented at the Annual Meeting of the Association of American Geographers.

"Digital Map Dependent Functions of Automatic Vehicle Location Systems", by C. B. Harris et al., IEEE Position and Location Symposium, pp. 79-87, 1988, IEEE CH2675-7.

"Digital Maps on Compact Discs", by H. J. G. M. Benning, Technical Paper Series 860125, Society of Automobile Engineers, 1986.

"Eva: An Electronic Traffic Pilot for Motorists", by O. Pilsak, Technical Papers Series 860346, Society of Automotive Engineers, 1986.

"Digital Map Data Bases for Autonomous Vehicle Navigation Systems", by E. P. Neukirchner et al.,

IEEE Position and Location Symposium, pp. 320-324, 1986, IEEE 86CH2365-5.

"Applications of the Compact Disc in Car Information and Navigation Systems", by M. L. G. Thoone et al., Technical Papers Series 840156, Society of Automotive Engineers, 1984.

"On Board Computer System for Navigation, Orientation, and Route Optimization", by P. Haeussermann, Technical Paper Series 840483, Society of Automotive Engineers, 1984.

"Electro Gyro-Cator: New Inertial Navigation System for Use in Automobiles", by K. Tagami et al., Technical Paper Series 830659, Soc. of Automotive Engineers, 1983.

"Navigation Systems Using gps for Vehicles", by T. Itoh, et al., Technical Paper Series 861360, Society of Automotive Engineers, 1986.

"Extending Low Cost Land Navigation Into Systems Information Distribution and Control", by S. K. Honey et al., IEEE Position and Locations Symposium, pp. 439-444, 1986, IEEE 86CH2365-5.

"Map Matching Augmented Dead Reckoning", by W. B. Zavoli et al., Proceedings of the 35th IEEE Vehicular Technology Conference, pp. 359-444, 1986, IEEE CH2308-5.

"Automated Provision of Navigation Assistance to Drivers", by Matthew McGranaghan et al., The American cartographer 14(2):121-138, 1987.

U.S. Patent

Jan. 5, 1993

Sheet 1 of 3

5,177,685

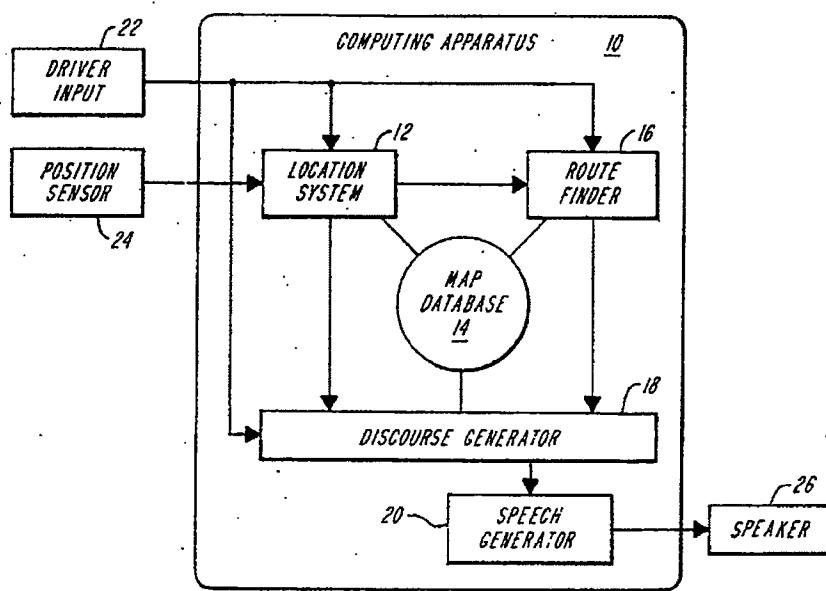


FIG. 1

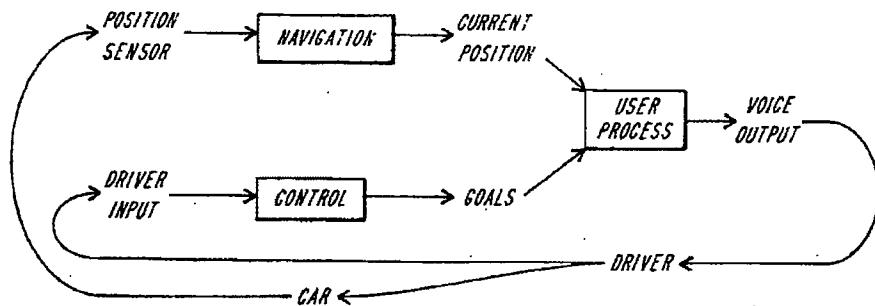


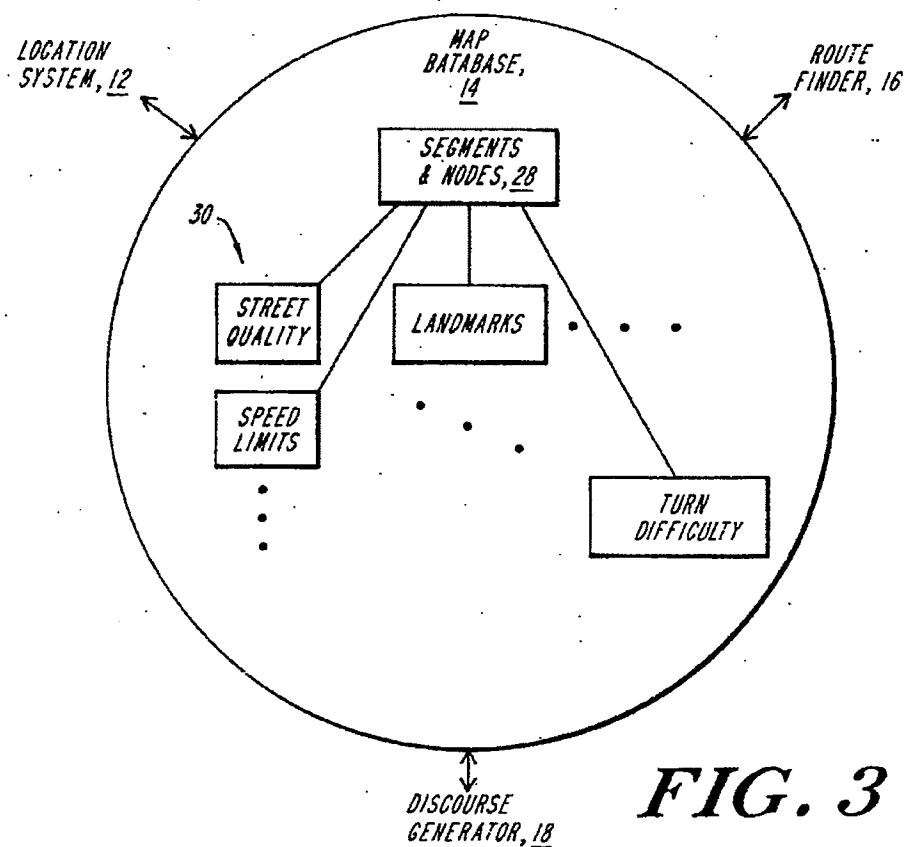
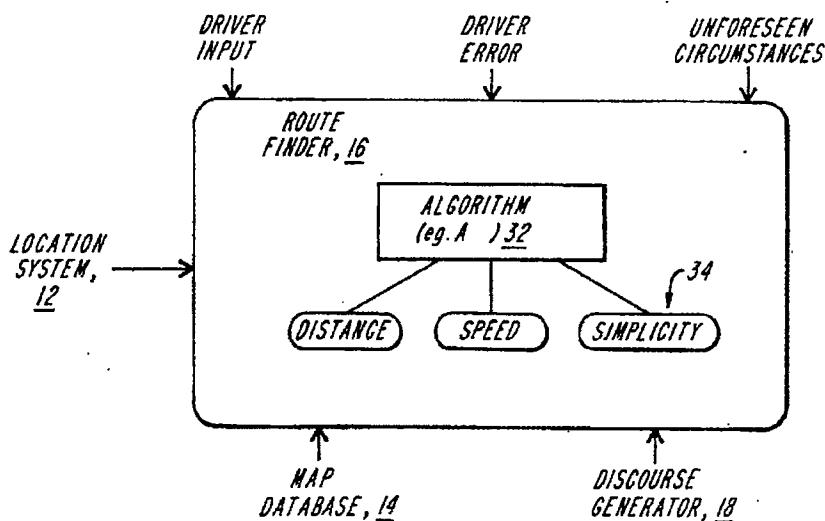
FIG. 2

U.S. Patent

Jan. 5, 1993

Sheet 2 of 3

5,177,685

***FIG. 3******FIG. 4***

U.S. Patent

Jan. 5, 1993

Sheet 3 of 3

5,177,685

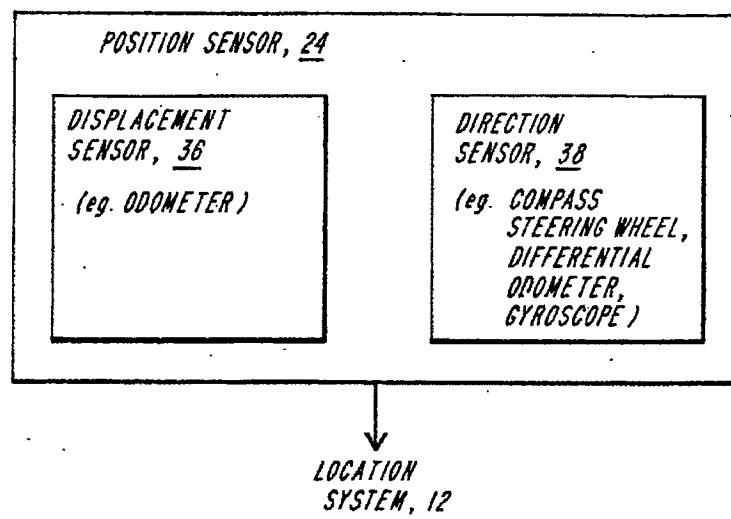


FIG. 5

5,177,685

1

2

AUTOMOBILE NAVIGATION SYSTEM USING REAL TIME SPOKEN DRIVING INSTRUCTIONS

BACKGROUND OF THE INVENTION

This invention relates to computerized automobile navigation systems, particularly to a system which calculates a route to a destination, tracks automobile location, and provides spoken instructions to the driver in real time as they are needed.

Navigation systems can be classified into three categories:

- Positioning systems tell you where you are.
- Orienting systems show the direction of your destination.
- Instructional systems tell you what to do to get to your destination.

A navigation system can provide one, two, or all of these services. Navigation systems can be further distinguished by how they provide the information:

- Verbal systems speak.
- Textual systems provide text.
- Graphic systems provide pictures.

Finally, systems can be classified as either real time or static. The categories of this classification are not independent. There can be no static positioning system, since one cannot predict the future position of an automobile.

There are several problems with static navigation systems. First, they do not help the driver follow the route. The driver must determine when to apply each instruction. A second problem is that since the instructions must be specified in advance, there is little to be done if the driver does not follow the instructions, which might happen from error, or because the instructions are wrong, or simply ill-advised (as when confronting a traffic jam).

Previous automobile navigation systems have used text or graphics to give navigation information. However, there are several disadvantages to presenting information visually. First, the driver must look at a display while driving, which makes driving less safe. For providing driving directions, visual displays are most easily used when they are least needed. Second, with respect to graphic displays, many people have difficulty using maps, making this mode of providing information undesirable. However, if speech is used, the driver's eyes are left free for driving. In addition, speech uses words, and can therefore refer to past and future actions and objects not yet seen. This is hard to do with symbolic displays or maps.

There is clearly a need for an instructional, verbal, real time automobile navigation system which can guide a driver to a destination much as a passenger familiar with the route would. The present invention meets that need.

SUMMARY OF THE INVENTION

The present invention, called the "Back Seat Driver", is a computer navigation system which gives spoken instructions to the driver of an automobile to guide the driver to a desired destination. Computing apparatus, installed either in the automobile or accessed through a cellular car phone, contains a map database and a route finding algorithm. A vehicle location system uses data from a position sensor installed in the automobile to track the location of the automobile. Discourse generating programs compose driving in-

structions and other messages which are communicated to the driver using voice generating apparatus as the driver proceeds along the route.

The important differences between The Back Seat Driver and other such systems are that the Back Seat Driver finds routes for the driver, instead of simply displaying position on a map, tells the driver how to follow the route, step by step, instead of just showing the route, and speaks its instructions, instead of displaying them. Each of these design goals has required new features in the programs or in the street map database.

The street map database of the Back Seat Driver distinguishes between physical connectivity (how pieces of pavement connect) and legal connectivity (whether one can legally drive onto a physically connected piece of pavement). Legal connectivity is essential for route finding, and physical connectivity for describing the route.

To find the fastest routes, the map database of the Back Seat Driver includes features that affect speed of travel, including street quality, speed limit, traffic lights and stop signs. To generate directions, the map includes landmarks such as traffic lights and buildings, and additional descriptive information about the street segments, including street type, number of lanes, turn restrictions, street quality, and speed limit. The map also preferably includes other features, such as time-dependent legal connectivity, and expected rate of travel along streets and across intersections. Positions are preferably stored in the map database in three dimensions, not two, with sufficient accuracy that the headings of the streets can be accurately determined from the map segments.

Driving instructions generated by the Back Seat Driver are modeled after those given by people. The two issues for spoken directions are what to say (content) and when to say it (timing). The content of the instructions tells the driver what to do and where to do it. The Back Seat Driver has a large taxonomy of intersection types, and chooses verbs to indicate the kind of intersection and the way of moving through it. The instructions refer to landmarks and timing to tell the driver when to act.

Timing is critical because speech is transient. The Back Seat Driver gives instructions just in time for the driver to take the required action, and thus the driver need not remember the instruction or exert effort looking for the place to act. The Back Seat Driver also gives instructions in advance, if time allows, and the driver may request additional instructions at any time. If the driver makes a mistake, the Back Seat Driver describes the mistake, without casting blame, then finds a new route from the current location.

Giving instructions for following a route requires breaking the route down into a sequence of driving acts, and knowing when an act is obvious to the driver and when it needs to be mentioned. This further requires knowledge about the individual driver, for what is obvious to one may not be so to another. The Back Seat Driver preferably stores knowledge of its users, and uses this knowledge to customize its instructions to the preferences of the users.

Speech, especially synthetic speech, as an output media imposes constraints on the interface. The transient nature of speech requires that utterances be repeatable on demand. The Back Seat Driver has the ability to construct a new utterance with the same intent, but not necessarily the same words, as a previous message.

5,177,685

3

Synthetic speech being sometimes hard to understand, the Back Seat Driver chooses its words to provide redundancy in its utterances.

An actual working prototype of the Back Seat Driver has been implemented. It has successfully guided drivers unfamiliar with Cambridge, Mass. to their destinations. It is easy to foresee a practical implementation in the future.

BRIEF DESCRIPTION OF THE DRAWING

FIG. 1 illustrates schematically the basic functional components of the Back Seat Driver in its preferred embodiment.

FIG. 2 illustrates the system processes of the preferred embodiment of the Back Seat Driver.

FIG. 3 is a schematic illustration of the map database.

FIG. 4 is a schematic illustration of the route finder.

FIG. 5 is a schematic illustration of the position sensor.

DESCRIPTION OF THE PREFERRED EMBODIMENT

The automobile navigation system according to the present invention is illustrated schematically in FIG. 1. The heart of the system is a computing apparatus 10 comprising a vehicle location system 12, a map database 14, a route finder 16, a discourse generator 18, and a speech generator 20. Driver input means 22 allows the driver to input to the computing apparatus 10 information such as a desired destination. A position sensor 24 measures automobile movement (magnitude and direction) and sends data to the location system 12 which tracks the position of the automobile on the map. The route finder 16 calculates a route to the destination. Based on the current position of the automobile and the route, the discourse generator 18 composes driving instructions and other messages according to a discourse model in real time as they are needed. The instructions and messages are sent to the speech generator 20 which conveys them to the driver by means of a speaker system 26. The speaker system may be that of the car's radio.

In FIG. 1, the computing apparatus is illustrated as a single entity. However, in other embodiments, the components may not all be implemented in the same piece of apparatus. For example, in one working prototype of the Back Seat Driver, the main computing apparatus is a Symbolics Lisp machine, but the location system is implemented separately by an NEC location system that tracks the position of the automobile using its own map database, and the speech generator is implemented separately by a Dectalk speech synthesizer. In another working prototype, the main computing apparatus is a Sun Sparc workstation. The map database for the Back Seat Driver can be provided on a CD-ROM, a floppy disk, or stored in solid-state memory, for example.

The components of the system and the system processes which coordinate their performance, particularly as embodied in the working prototypes, are discussed in the sections which follow. Aspects of the invention are also described in the following sources, which are hereby incorporated by reference:

1. "Synthetic speech for real time direction-giving," by C. M. Schmandt and J. R. Davis, *Digest of Technical Papers, International Conference on Consumer Electronics*, Rosemont, Ill., Jun. 6-9, 1989.
2. "Synthetic speech for real time direction-giving," by C. M. Schmandt and J. R. Davis, *IEEE Transactions on Consumer Electronics*, 35(3): 649-653, August 1989.

4

on Consumer Electronics, 35(3): 649-653, August 1989.

3. "The Back Seat Driver: Real time spoken driving instructions," by J. R. Davis and C. M. Schmandt, *Proceedings of the IEEE Vehicle Navigation and Information Systems Conference*, Toronto, Canada, September 1989.
4. "Back Seat Driver: Voice assisted automobile navigation," by J. R. Davis, Ph.D. thesis, Massachusetts Institute of Technology, September 1989.

MAP DATABASE

The map database for the Back Seat Driver in the working prototypes originated as a DIME (Dual Independent Map Encoding) file, a map format invented by the U.S. Census Bureau for the 1980 census. Implementing the Back Seat Driver required extending the DIME map format in a number of ways to make it useful for route finding and route describing.

The basic unit of the DIME file is the segment. A segment is a portion of a street (or other linear feature such as a railroad, property line, or shoreline) chosen to be small enough that it is a straight line and has no intersection with any other segment except at its endpoints.

The two endpoints of a segment are designated FROM and TO. If the segment is a street segment (as opposed to, say, a railroad) and has addresses on it, then the FROM endpoint is the one with the lowest address. Otherwise, the endpoint labels are chosen arbitrarily. A segment has two sides, left and right. The sides are chosen with respect to travel from the FROM endpoint to the TO endpoint. A navigator using a DIME file can find the location of an address along the segment by interpolating the addresses between the low and high addresses for the two endpoints. The DIME file is suited to determining the approximate position of a building from its street address.

Attributes of a DIME file segment include: its name (40 characters), its type (a one to four character abbreviation such as "ST"), the ZIP code for each side, and the addresses for each endpoint and each side. At each endpoint of a segment is a pointer to a node. A node represents the coordinates of that endpoint and the set of other segments which are physically connected at that endpoint. Segments share nodes. If any two segments have an endpoint at the same coordinate, they will both use the same node for that endpoint.

A vehicle navigation system using a DIME file can represent the position of a vehicle on the map by a structure called a position. A position has three parts: a segment, an orientation, and a distance. The segment is one of the segments from the map database, the orientation specifies the direction the vehicle is travelling (towards the TO or FROM endpoint), and the distance is the distance from the FROM endpoint of the segment, no matter which way the vehicle is oriented. When travelling towards the TO endpoint of the segment, distance increases, when travelling towards the FROM endpoint, it decreases.

The DIME file is not adequate for routing finding and is only marginal for generating route descriptions. The most important problem with the DIME format is that it indicates only if two segments are physically connected (that is, if they touch), but not whether they are legally connected (i.e. whether it is legal to travel from one to the other). Legal connectivity is crucial for route finding. However, legal connectivity does not

5,177,685

5

replace physical connectivity; route description requires information about physical connections as well. Physical connectivity also affects route finding directly when seeking the simplest route, since ease of description is determined in part by physical connectivity.

The most significant extension of the DIME file format required for its use in a vehicle navigation system is the explicit representation of legal connectivity. This can be accomplished by adding a legal connection list at each endpoint of a segment to indicate all segments which are legally accessible from that endpoint. This allows the route finder to consider only legal paths. To the inventor's knowledge, this has not been included in any other navigation system.

Another problem with the DIME file is that it is a planar graph. This means that no two segments can cross except at an intersection, so there is no way to correctly represent an overpass, for example. The DIME format represents an overpass by breaking both streets at the point where they cross, and creating a fictitious intersection even though the segments do not touch in reality. These false intersections are particularly troublesome since DIME does not represent legal connectivity, so it appears possible and legal for a car to jump straight up and turn onto the overpass.

Points in the map database for a vehicle navigation system are therefore preferably three-dimensional. Route descriptions then provide better knowledge of the underlying topography. Stopping distance is affected by slope, so instructions must be given sooner when traveling down a hill. Slope affects safety. The route finder should avoid steep slopes in snowy weather. Finally, distance between points will be more accurate when change in altitude is considered. Roads designed for high speed may be more level than the underlying topography. They may be elevated or they may be depressed. A road which is not at grade will not have the slope of the land beneath it.

Coordinates in the DIME file are stored in ten thousandths of a degree. This means that the position of an endpoint in the map differs from the true position by as much as 6.5 meters in latitude and 5 meters in longitude at the latitude of Boston. This inherent position error causes problems because it introduces error in length and in heading. Uncertainty in heading causes uncertainty in the angle between two segments. A straight street can appear to wobble if it is made of many short segments. Segment "wobble" causes problems for a route finder, makes it hard to generate correct descriptions, and interferes with position determination.

DIME file segment "wobble" can be corrected for by assuming that the angle between two streets is the smallest possible value. However, this sometimes overestimates the speed an intersection can be travelled through. Uncertainty in the angle of segments at an intersection also makes it difficult to describe the intersection correctly and interferes with navigation because it makes it difficult to compare compass headings with the heading of a street.

A richer taxonomy of street types than that provided by DIME is preferable for a vehicle navigation system. Important categories of streets are: ordinary street, rotary, access ramp, underpass, tunnel, and bridge. Preferably, non-streets such as railroad, water, alley and walkway are also included.

The DIME file records a small amount of information about each segment. For a vehicle navigation system, additional attributes are preferably added to make bet-

6

ter descriptions. Important additional attributes are street quality, divided roads, signs, traffic lights, stop signs, buildings, other landmarks, lane information, and speed limit.

The street quality can be, for example, a number from 1 ("super") to 4 ("bad") which combines the ease of locating and following the street and the expected rate of travel along it. The street quality attribute should be used by both the route finder and the route describer.

The identification of divided roads is necessary to avoid U Turns where they are not possible, although it is preferable to make U Turns only if there is no other alternative. In addition, the route finder should recognize that a divided road is safer than an undivided road.

Sign and exit numbers are preferably stored in the map database as connection cues, which are text strings that give cues for moving from one segment to another. Every cue has a type which tells the kind of cue, e.g. sign or exit-number. There may be more than one connection cue for a given pair of segments, but there should never be more than one of a type.

The most useful landmarks are traffic lights. Traffic lights are preferably stored independently for each endpoint of each segment, since the presence of a light at one segment of an intersection does not imply that all other segments at the intersection have a light.

Two types of buildings which are especially useful as landmarks are toll booths and gas stations. Toll booths can be stored as connection cues. Gas stations can be stored in the services database described below. However, a preferred approach is to index gas stations (and other buildings) by street.

Roads often have more than one lane. Selecting the proper lane can make travel faster, and it may even be mandatory, since certain turns may only be possible from some lanes. The map database therefore preferably contains the number of lanes for both directions on a segment, and whether one or more lanes is reserved for turn restrictions.

The map database also preferably includes time dependent legal connectivity. Sometimes a given turn will be prohibited at certain hours of the day, typically rush hour. Additionally, lanes sometimes switch direction during the day to accommodate rush hour traffic, and some lanes are reserved for carpools during rush hour.

The expected rate of travel is not necessarily a function of street quality. Although there is a correlation, travel rate is preferably a separate segment attribute. One reason is that travel rate, unlike quality, changes during the day. A model of traffic flow like that of an experienced driver (i.e. it should know what "rush hour" means) is preferably implemented in the map database.

Some turns, though legal, are difficult to make. The route finder preferably avoids these turns if possible. To an extent, the difficulty of a turn is implicit in the quality of the participating street segments, but an explicit model in the map database is preferred.

Some lanes or streets are restricted to certain kinds of traffic (car pools, no commercial vehicles). Also important are height restrictions, since some underpasses are so low that tall vehicles will not fit under them. This information is preferably included in the map database.

At some lights it is permitted to make a right turn at a red light after a full stop. Right turns here will be no slower than rights turns at a stop sign, so the route finder should prefer such intersections to those that do not permit it. Also, traffic lights have differing cycle

5,177,685

7

lengths. The map database preferably includes this information.

Local knowledge is also preferably included in the map database. These are facts about how people and institutions act on or near the road; e.g. that a speed trap is here, or that this road is one of the first ones plowed after a snow storm.

The Back Seat Driver should allow the driver to select famous destinations by name in addition to address by including this information in a database, and this database should be integrated with the services database, discussed below. The Back Seat Driver should also support names of buildings and office plazas made up by developers without reference to the street names.

Service locations are preferably stored in a services database. This database lists services such as gas stations, automatic teller machines and stores. For each service is recorded the name of the establishment, the address, phone number, and hours of operation. This allows the Back Seat Driver to select the closest provider of a service known to be open. The database can also be used as a source of landmarks when giving directions.

The map database preferably contains information on the division of the city into neighborhoods. This is useful for selecting an address. The postal ZIP code is not good for classifying neighborhoods.

Pronunciation information is preferably stored in a database for those place names which are easily mispronounced by the speech synthesizer. It would also be desirable to record which of those names have unusual spellings. This would allow the system to warn the driver to be alert for signs that might otherwise surprise her. Note that the driver only hears the name of a street. and has to guess how it is spelled from the sound she hears.

Abbreviations are preferably included to allow the user to enter certain street names in abbreviated form. A second use for abbreviations is to supply alternate spellings for streets, for example, to allow the driver to spell "Mt Auburn" as "Mount Auburn".

An almanac is preferably included to list the time of sunrise and sunset for the city. Arrangements can be made to either purchase this database or locate a program which calculates it, for arbitrary position and date.

A problem for a practical Back Seat Driver is how to keep the map database accurate, since the streets network is constantly changing. Over time, new streets are constructed, old streets are renamed or closed. These kinds of changes are predictable, slow, and long lasting. Other changes are unpredictable, quick, and transient. A road may be closed for repairs for the day, blocked by a fallen tree, or full of snow. Such changes are usually short lived. Thus, the Back Seat Driver needs the ability to change legal connectivity dynamically. In addition, the route finder should preferably have the ability to avoid congested roads caused by rush hour or accidents, for example. The map database is therefore preferably continuously updated by some form of radio broadcast by an agency that monitors construction and real time traffic conditions.

The Census Bureau, in cooperation with the United States Geological Survey, has designed a new map format known as TIGER (Topologically Integrated Geographic Encoding and Referencing) which has several improvements over the DIME format, but

which is still a planar graph representing only physical connectivity. The map database for a Back Seat Driver could be also be originated from a TIGER file as long as the extensions discussed above were implemented.

The map database is shown schematically in FIG. 3. In the preferred embodiment, the map database 14 includes, as its basis, a file 28 of segments and nodes. File 28 may be an original file or may be adapted from a DIME file or a TIGER file by adding the above-described extensions. In addition, the map database 14 may include optional features 30, as described above.

ROUTE FINDER

Finding a route between two points in a street network is an example of searching a general graph. The task is to find a sequence of segments that lead from the origin to the destination. There are usually a great many distinct ways of getting from one place in the city to another, some better than others. Graph search algorithms differ in the quality of the solution they find and the time they require. The Back Seat Driver requires an algorithm that finds a good route in a short time.

The route finder of the working prototypes of the Back Seat Driver is based on an A* search algorithm. The A* algorithm is a form of best-first search, which itself is a form of breadth-first search. These searching techniques are well-known and are described in detail in Davis, 1989, cited above.

In a breadth-first search, a tree of all possible decisions is divided into levels, where the first level actions are those leading from the root, the second level actions are those that come from situations after first level actions, and so on. All actions at a given level are considered before any at the next higher level. While the breadth-first search is operating, it maintains a list of all possible partial routes and systematically examines every possible path from the end of every partial route to compile a new list of partial routes. This search procedure finds the path with the fewest segments. However, this is not necessarily the best path. To be sure of finding the best path, the search cannot stop when the first path is found, but must continue, expanding each path, until all paths are complete. This is not at all desirable, since there could be (and in fact will be) many paths.

The best-first algorithm solves this problem by keeping track of the (partial) cost of each path, and examining the one with the smallest cost so far. This requires a function that can compare two routes and produce a numeric rating. Such a function is called a metric. To further reduce the cost of searching, before adding a segment to a path, the best-first search checks to see whether it is a member of any other path. If it is, it is not added, for presence on the other path means that the other path was a less expensive way of reaching the same segment.

Best-first search finds the best solution and requires less time than exhaustive breadth-first search, but it still must consider partial solutions with an initial low cost which prove expensive when complete. The A* algorithm avoids wasting time on such falsely promising solutions by including an estimate for the completed cost when selecting the next partial solution to work on. The cost estimate function is $f^*(r) = g^*(r) + h^*(r)$, where r is a route, $g^*(r)$ is the known cost of the partial route, and $h^*(r)$ is the estimate of the cost to go from the endpoint of the route to the goal. The h^* function must have the property of being always non-negative and

5,177,685

9

never over-estimating the remaining cost. An h^* meeting these two conditions is said to be admissible. It should be obvious that if h^* is chosen to be always zero, then A* search is just best-first search. In applying A* to finding routes on a map, h^* is just the cartesian distance between the endpoint of the partial route and the destination point. It is certain that no route will be shorter than the straight line, so this estimate is never an over estimate. A* search is more efficient than best-first.

The A* algorithm finds the optimum route, but the Back Seat Driver might be better served with an algorithm that finds a reasonable route in less time. This is especially true when the vehicle is in motion. The longer the route finder takes, the greater the distance that must be reserved for route finding. As this distance becomes larger, it becomes harder to predict the future position of the car. This can be done by choosing an h^* which multiplies the estimated distance remaining by a constant D. Setting D greater than one makes h^* no longer admissible, since the estimate might exceed the actual cost by a factor of D. The resulting routes are no longer optimal, but are still pretty good. The effect is to make the algorithm reluctant to consider routes which initially lead away from the goal.

The route finder preferably uses a value of 2 for D. This yields the greatest increase in payoff. A possible improvement is to run the route finder twice, first with a high value of D to find an initial route in order to begin the trip, and then with a low D to search for a better route, using spare time while driving.

Preferably, three different metrics are used. The distance metric finds the shortest route, the speed metric finds the fastest route, and the ease metric finds the easiest route. The metric for distance is just the sum of the lengths of the component segments. The other two metrics are more complicated than the distance metric, because they must consider intersections as well as segments. In general there is a cost to travel along a segment and a cost to get from one segment to another. All costs are expressed as an "equivalent distance" which is the extra distance one would travel to avoid the cost.

The metric for speed estimates the cost for traveling along a segment by multiplying its length by a constant which depends upon the quality of the street. In principle, one could calculate expected time by dividing length by the average speed on the segment were this quantity available in the database. Examples of appropriate constants are:

Quality	Factor
super	1
good	1.2
average	1.5
bad	2.0

All multiplicative constants must be greater than or equal to one to ensure that the cost of a route is never less than the straight line distance between two points. This condition is essential for the correct operation of the A* search algorithm, since the estimation function (g^*) must always return an under-estimate.

The time to cross an intersection is preferably modeled by a mileage penalty which depends upon the nature of the intersection. Examples of appropriate penalties are:

10

Factor	Cost	Reason
turn	1/2 mile	Must slow down to turn
left turn	1/2 mile	May have to wait for turn across traffic flow
traffic light	1 mile	Might be red

If the segment is one-way, the penalties should be cut in half, since there will be no opposing traffic flow. The turning penalties should be computed based only on the angle between two segments, not on the segment type or quality.

The metric for ease seeks to minimize the driver's effort in following the route. Again, driver's effort is the sum of the effort to travel along a segment and the effort to get from one segment to another. Travel along a segment depends upon its quality. Turns of every sort should be penalized equally, since they all require decisions. The intention of this metric is to find routes which require the least amount of speaking by the Back Seat Driver, leaving the driver free to concentrate on other matters.

If the driver leaves the route, the Back Seat Driver must immediately inform the driver and begin to plan a new route. Route planning after a mistake is no different from any other time, except that the vehicle is more likely to be moving. In the working prototypes, when the car is moving, the Back Seat Driver first estimates the distance the car will travel during the route finding process by multiplying the current velocity by the estimated time to find the route. Then it finds the position the driver will reach after traveling this distance, assuming that the driver will not make any turns without being told to do so. It then finds a route from this extrapolated position to the goal. Finally, it finds a route from the car's actual position to the estimated starting position. This second route is so short that the car is unlikely to move far during the time it is computed.

The route finder of the working prototypes estimates the time to find the route between two points by multiplying the distance between them by a constant. This constant was initially determined by running the route finder for 20 randomly selected pairs of origins and destinations. As the Back Seat Driver runs, it accumulates additional values for the constant.

A problem is how to reliably detect when the driver has left the route. With the extended DIME format of the working prototypes, if the driver turns into a gas station, for example, the system will believe, falsely, that the driver has turned onto some street, because the street map includes only streets, and not other paved areas such as parking lots and filling stations. From this false belief, the system will conclude that the driver has made a mistake. However, this problem can be solved by increasing the detail of the map.

Sometimes the driver will choose to not follow a route for good reasons that the Back Seat Driver is unaware of, perhaps because the road is blocked or because of a traffic jam. For the first case, the driver should be provided an "I Can't Do It" button or other means to inform the Back Seat Driver that the road is (temporarily) blocked. Once informed, the Back Seat Driver must automatically find a new route. For the second case, the driver's only recourse is to cancel the current trip (by pushing another button, for example), and, once out of the situation, re-request a route to the original destination. It is essential, though, that the

5,177,685

11

driver either notify the Back Seat Driver of the impossibility of the requested action or cancel the trip, because otherwise the Back Seat Driver will treat the deviation from the route as a mistake, and continue to attempt to find a new route, which may very well lead back through the street the driver is trying to avoid.

The route finder is shown schematically in FIG. 4. In the preferred embodiment, the route finder 16 includes, as its basis, an algorithm 32. Algorithm 32 may be, for example, an original algorithm based on a best-first search algorithm the A* algorithm, or a modified A* algorithm. In preferred embodiments, the route finder is adapted to find the best route according to any one of three cost metrics 34: distance, speed, simplicity. The route finder calculates a new route in the event of driver error or unforeseen circumstances, as indicated.

LOCATION SYSTEM AND POSITION SENSOR

The Back Seat Driver must know the position of the vehicle. This can be achieved using available technology adapted for the requirements of the Back Seat Driver. At a minimum, the location system for a vehicle navigation system must determine the vehicle position to the nearest block. If it is to tell the driver when to turn, it must be able to distinguish between the closest of two adjacent turns.

Consideration of the Boston street map shows that it has many streets which are both short and a possible choice point. Based on a study of the percentage of segments which are shorter than a given length, an accuracy of 10 meters is desirable. This is a higher accuracy than has been specified in prior art approaches (see Davis, 1989, cited above). The Back Seat Driver's use of speech imposes strict requirements on position because of limitations on time. Unlike a display, speech is transient. An action described too soon may be forgotten. The Back Seat Driver is intended to speak at the latest time that still permits the driver to act. Allowing two seconds for speech, a car at 30 mph covers 27 meters. This consideration suggests a minimum accuracy of 15 meters.

Location systems can be divided into two categories: Position finding systems determine position directly by detecting an external signal.

Position keeping (dead reckoning) systems estimate the current position from knowledge of an earlier position and the change in position since that position.

All existing position finding systems use radio signals. The broadcast stations may be located on street corners, seacoasts, or in orbit around the earth. Systems differ in their range, accuracy, and cost. A survey of those systems which might plausibly be used for automobile navigation is included in Davis, 1989, cited above.

Position keeping (dead reckoning) systems obtain position indirectly, by keeping track of the displacement from an originally known position. One can measure displacement directly, or measure velocity or acceleration and integrate over time to obtain displacement.

The forward motion of a car is measured by the odometer. On late model cars, the odometer cable has been standardized. It revolves once every 1.56 meters. This is a reliable measure of forward progress, as long as the wheels do not slip. Measuring direction, though, is more difficult. Among the possibilities are:

magnetic compass A magnetic compass has the advantages of small size and ease of use, but is unreliable because of variation between magnetic and true north

12

and deviation due to the ferrous material of the car and magnetic flux arising from electric currents within the car.

steering wheel The steering wheel could be instrumented to measure the amount of turning.

differential odometer When a car turns, the two rear wheels travel different distances, and thus rotate at different rates. Measuring the difference in rotation provides an indication of amount of turning. This differential rate of rotation is just what is measured by anti-skid brakes, so no additional instrumentation is required to obtain this measure for an automobile suitably equipped.

gyroscope Gyroscopes measure angular acceleration. The familiar rotation gyroscope and esoteric laser ring gyroscope are not suitable for automotive use because they are too expensive. Lower cost alternatives are the rate gyro and the gas jet gyro. The rate gyro measures angular acceleration in a vibrating piezo-electric substance. The gas gyro measures turn (or yaw) rate. In this design, a jet of gas travels down the center of a sealed tube. Anemometers are placed on either side of stream. When the car turns, the stream is deflected and the velocity is measured. The velocity of the gas at the anemometer is proportional to the turn rate. Such devices can measure turn rates of up to 100 degrees per second, with a noise of about one half degree/second.

The position sensor is shown schematically in FIG. 5. As indicated, it includes a displacement sensor 36 and a direction sensor 38.

A position keeping system with no error could be calibrated when installed, and then maintain its own position indefinitely. Unfortunately, errors arise in measuring both distance and heading. Sources for error include difference in tire pressure, composition and wear, slipping, cross steering from winds, change in tire contact path in turns, magnetic anomalies, and gyro noise. The NEC dead reckoning system, employed in the prototypes of the Back Seat Driver, accumulates about one meter of error for every ten meters traveled. The error is even worse when traveling near railroads because the NEC system uses a magnetic compass.

Some dead reckoning systems recalibrate themselves to eliminate systematic errors. Such recalibration is possible when the vehicle is at a known position or when stopped. One way to correct dead reckoning errors is to use knowledge of the road network as a constraint on position, in what is known as map matching. Map matching requires that the position keeping system have a map of the locale where the vehicle is being driven, and is based on the assumption that the vehicle is always on a street present in the map. If the estimated position falls to one side of the road, the estimate can be corrected. When the vehicle makes a turn, the system assumes the vehicle is at the closest intersection, and thus the absolute position can be estimated. Every dead reckoning system uses some form of map matching. Map matching reduces the stringency of position keeping, but accuracy remains a concern, since the system must maintain its position when the driver drives off the map, e.g. into a driveway or a parking lot.

In the working prototypes, a unit built by NEC Home Electronics, Ltd. is employed. It is a dead-reckoning position keeping system which uses speed and direction sensors. To compensate for error, it uses map matching on a map database stored on CD-ROM. The unit is described in "CD-ROM Assisted Navigation

13

Systems" by O. Ono, H. Ooe, and M. Sakamoto, in *Digest of Technical Papers, IEEE International Conference on Consumer Electronics*, Rosemont, Ill., Jun. 8-10, 1988.

As implemented in the working prototypes, the map database used by the location system is completely distinct from the map database used by the route finder and discourse generator. This is unfortunate since the maps will not always agree unless they are kept equally up-to-date. However, in other embodiments within the scope of the invention, the location system uses the computing resources and map database of the main computing apparatus illustrated in FIG. 1. Positioning systems for the Back Seat Driver preferably combine position keeping and position finding, since neither alone will work all the time. A position keeping system needs periodic corrections, but a position finding system that depends on radio reception will not work in tunnels or bridges. Hybrid systems which could be used by the Back Seat Driver are referenced and discussed in Davis, 1989, cited above.

DISCOURSE GENERATOR

The Back Seat Driver attempts to provide instructions to the driver as a passenger in the car familiar with the route would. The content and timing of the instructions and other messages described below are based on a study of natural driving instruction described in detail in Davis, 1989, cited above.

To the Back Seat Driver, a route is a sequence of street segments leading from the origin to the destination. Each connection from one segment to another is considered an intersection, even if there is only one next segment at the intersection. At any moment, the car will be on one of the segments of the route, approaching an intersection. The task of the Back Seat Driver is to say whatever is necessary to get the driver to go from the current segment, across the intersection, to the next segment of the route. Most often, nothing need be said. But at other times, the Back Seat Driver will need to give an instruction.

Instructions must use terms familiar to the driver. An example is what to say at a fork in the road. Considering only topology, there is no difference between a fork and a turn, but it would be confusing to call a fork a turn.

The two key issues in describing a route are deciding what to say and deciding when to say it. There is a tradeoff between these two factors. At one extreme are directions given completely in advance, with no control over when the driver reads them. A direction of this kind might be: "Go half a mile, then take a left onto Mulberry Street". A driver following such an instruction must use the odometer to estimate distance or look for a street sign. The instruction itself does not say when to act. The other extreme are instructions which rely totally on timing for success. Such an instruction might be: "Turn left now".

An intersection type is called an act because the important thing about an intersection is what action the driver takes to get across it. The Back Seat Driver is preferably implemented with an object-oriented programming methodology, so for each act there is an expert (an object) capable of recognizing and describing the act. The Back Seat Driver generates speech by consulting these experts. At any moment, there will be exactly one expert in charge of telling the driver what to do. To select this expert, the Back Seat Driver asks each expert in turn to decide whether it applies to the

14

intersection. The experts are consulted in a fixed order, the most specific ones first. The first expert to claim responsibility is selected. This expert then has the responsibility of deciding what (if anything) to say.

Each act has a recognition predicate which is used to determine if a given intersection should be classified as that act. A predicate can consider topology, geometry, the types of street involved, or any other factor. The predicate also decides whether the move is obvious, that is, the driver can be trusted to do it without being explicitly told to do so. Actions that are obvious are not described. If the next action is obvious, the Back Seat Driver looks ahead along the route until it finds one which is not obvious. There will always be at least one, because stopping at the end is never obvious.

The acts in the working prototypes include CONTINUE, FORCED-TURN, U-TURN, ENTER, EXIT, ONTO-ROTARY, EXIT-ROTARY, STAY-ON-ROTARY, FORK, TURN and STOP.

A CONTINUE is recognized when the driver is to stay on the "same" road. Almost always, a continue is obvious and nothing should be said. The continuation of a street depends on the type of street: from a rotary, it is the next rotary segment; from an access ramp, if there is exactly one next segment, that is the continuation, otherwise there is no obvious next segment; otherwise, it is the one segment that requires no more than 30 degrees of angle change (if there is exactly one, and if it is not a rotary) or the one segment with the same name (if there is exactly one). The reason for comparing names is not because the driver is aware of the name, but because the designer who named the street was. The assumption is that if two segments have the same name, they are the same street, and that is why they have the same name. This "sameness" is presumably reflected in details not captured by the map, for example continuity of painted centerline. There are many places in the area where the obvious "straight" continuation of a segment is at an angle as great as 45 degrees, but it would not be right to call this a turn.

A FORCED-TURN is an intersection where there is only one next street segment where the road bends more than 10 degrees. Even though there is no decision to make at a forced turn, it is useful to mention because it strengthens the driver's sense that the Back Seat Driver really knows about the road conditions. A forced turn is not worth mentioning if both segments are part of a bridge, a tunnel, or an access ramp, or if the angle is less than 20 degrees.

The U-TURN action is recognized when the heading of the car is the opposite of what it should be. Recall that a route is a sequence of segments and endpoints. At all times the car will be on one of the segments in the sequence. If the car's orientation is not the same as the endpoint in the path, then the driver must turn around. Preferably, the route finder only calls for a U Turn if there is no other way.

To ENTER is to move onto a super street (or an access ramp that leads eventually to a super street) from an ordinary street, but not from a super street or an earlier access ramp. Similarly, to EXIT is to move from a super street onto a street with lesser quality that is either an access ramp or has a different name. Some super streets are not uniformly super and it would not be right to call the change in quality an exit.

To go ONTO-ROTARY, to STAY-ON-ROTARY, and to EXIT-ROTARY are acts which can be correctly

5,177,685

15

described only if the street map database includes an explicit marking of streets as rotaries.

At a FORK, there must be at least two alternatives, all within a narrow angle, and none of the branches must be the obvious next segment—that is, the branches must all be more or less equal. Either all the alternatives must be access ramps, or none of them must be. A branch can only be considered obvious if it is the only branch with the same level of quality, or if it is markedly straighter than the others, or if it is the only one with the same number of lanes, provided that all of these clues agree. If one branch is stronger than the others, the intersection is not a fork. It is either a continue or a turn.

The STOP action is recognized when the vehicle is on the destination segment. Finally, a TURN is an intersection not handled by one of the above cases. The greatest weakness of the above approach is that the recognition predicates are sensitive to small changes in the angles between segments. It is not likely that people use absolute numbers (e.g. 10 degrees) as cut-off values for their determination of how to describe an intersection. More likely, different classifications compete. Still more important, people making classifications use visual cues, not just facts from the map.

Each act has a description function to generate a description of the action. The description function takes inputs specifying the size of the description (brief or long), the tense (past, present or future), and the reference position. A short description is the minimum necessary for the act. It is typically an imperative (e.g. "Bear left."). A long description includes other facts about the action, an expression indicating the distance or time until the act is to be performed, and possibly information about the next act, if it is close. The reference position is a position (along the route) from which the action is to be described.

A brief description consists only of a verb phrase. The verb depends on the type of act and perhaps on the specifics of the act. Besides the verb itself, the verb phrase must say which way to go. In most cases, the word "left" or "right" is sufficient. Where it is not, the possibilities are to use a landmark or to describe the turn. A landmark can be either in the appropriate direction ("towards the underpass") or the other direction ("away from the river"). Specifying direction with a landmark has the advantage that some drivers confuse left and right sides, or mishear the words, so it is a redundant cue. Also, it increases the driver's confidence that the system really knows what the land looks like. A description of the turn can mention either quality or the relative angle of the desired road. The angle must be described qualitatively (more or less "sharp"). It would be more precise to use the angular distance (e.g. "turn right 83 degrees"), but drivers would not understand it. Preferably, the expert for each act follows a protocol which includes:

recognize?—is a proposed movement an example of this kind of driving act?

instruction-vp—generate a verb phrase describing this act

instruction-np—generate a noun phrase describing the act

position-to-doit—the position where the driver would begin carrying out the act

obvious?—would the driver do this act without being told?

16

sentences—generate all sentences needed to describe this act

congratulate?—should the driver be congratulated after carrying out this kind of act

The following sample is a Back Seat Driver description of the left turn from Fulkerson Street to Main Street in Kendall Square, Cambridge, Mass.:

Get in the left lane because you're going to take a left at the next set of lights. It's a complicated intersection because there are two streets on the left. You want the sharper of the two. It's also the better of them. After the turn, get into the right lane.

This instruction begins with a piece of lane advice, an action to be taken immediately, then describes an action in the near future. The action is a turn, though that word is not used explicitly. It tells the direction of the turn (left) and specifies a landmark (the lights) that says where the turn is. In many cases, this would be enough, but here there are two streets on the left, so the instruction goes on to specify the desired road in two ways (by comparative position and relative quality). Finally, it concludes with some lane advice to be executed during (or just after) the act.

The above example is the most complicated text that the Back Seat Driver prototypes have produced. Length and detail are not virtues in giving directions. The Back Seat Driver produces a text this long only because it does not have better means to make the driver follow the route. If a shorter text would accomplish the same aim, it would be better.

Besides telling drivers what to do, the Back Seat Driver must also tell them when to do it. One way to do this is by speaking at the moment to act, but it is useful to also give instructions before the act, if time permits. This allows time for preparation, if required, permits the driver to hear the instruction twice, and also spares the driver the need to be constantly alert for a command which must be obeyed at once.

When an act is more than a few seconds in the future, the Back Seat Driver uses a long description, which includes one or more cues which either describe the place for the act, the features of the road between the current location and the place, or the distance or time until the act. This description should be so clear that the driver cannot only recognize the place when it comes, but can also be confident in advance that she will be able to recognize the place.

The Back Seat Driver preferably uses a landmark as a cue when it can. A numeric distance is the cue of last resort. However, some drivers prefer to also hear distances, especially if the distance exceeds a certain threshold. Therefore, a parameter is preferably included in the user-model, described below, for this minimum distance expressed as a number. If the distance is below this, a qualitative phrase is produced by the discourse generator, if above, a number is produced. The cutoff can be zero, in which case numbers are always used, or set to an infinite value, in which case they never are.

A cue is expressed either as a full sentence ("Drive to the end of the street, then . . . ") or a proposed preposition phrase ("At the next set of lights, . . . "). Research has shown that a cue should not be expressed by a preposition after the verb as in "Take a left at the lights," because some drivers start to take the left as soon as they hear the word "left". This may be because syn-

5,177,685

17

thetic speech does not provide enough intonational cues for the driver to reliably predict the length of the sentence, leading the driver to act on syntactic information alone, and thus taking the sentence to be complete as soon as the word "left" is heard.

The description of a road feature depends upon whether or not it is visible. If it is, it can be referred to with a definite article ("the rotary", "the overpass"). If not, an indefinite article is used. The program cannot tell whether an entity is actually visible, so it uses distance as an approximation. If the feature is closer than one tenth of a mile, it is considered to be visible.

A special case of cues is when the driver is at the place to act. When stopped a few meters from the intersection, it is wrong to say "Turn at the next lights" even if it is literally true. In the working prototypes, the Back Seat Driver thinks of itself as being at that intersection if its less than thirty yards away, except that if there is a stop light at the intersection and the car is not moving, then the intersection distance is fifty yards, since cars might be backed up at such an intersection. When at an intersection, the Back Seat Driver should say "Take a left here" rather than "Take a left now" because drivers waiting for a traffic light will rightly resent being told to do something they have good reason not to do.

Traffic lights are very good landmarks because they are designed to be easily seen and drivers have an independent reason to watch for them, namely a desire to avoid accidents. When referring to a traffic light, if the car is at the intersection for the lights, the Back Seat Driver should use a proximal deictic ("this" or "these", as opposed to the distal "that" or "those") to show it means the lights that are here.

The Back Seat Driver preferably has a database of buildings as part of its directory of services. If it finds a building on the corner, it should include it as a potential landmark: e.g. "Look for Merit Gas on the left side".

Other landmarks are features of the road, such as underpasses, bridges, tunnels, bends in the road, and railroad crossings. Still another potential landmark is the road coming to an end. This is a landmark that is impossible to miss. However, research has shown that if the Back Seat Driver says "Drive all the way to the end, then . . .", some drivers take "the end" to mean not "the farthest you can go along this road" but just "the next intersection".

A street name can be a landmark, but not a good one, unless the driver already knows the street. There are several reasons why street names should not be used. First, the driver may not hear the name correctly. Second, the driver may hear the name, but not know how to spell the name after hearing it, so she may not recognize the name in its printed form. This is especially a problem when the driver is from out of town. Finally, even if the driver knows the spelling, street signs are often missing, turned around, or invisible due to weather or darkness. Despite all the problems that come with using street names, many drivers ask for them. To accommodate them, a parameter in the user-model is preferably included to control the use of names. If set, names are supplied as part of the instruction. When names are included, they are preferably attached at the end of the instruction ("Take the second left. It's Elm Street.") rather than directly ("Take the second left onto Elm Street."), which weakens their salience somewhat, and makes them more of a confirmatory cue than an essential one.

18

Signs can be important landmarks. A problem with using signs as cues occurs, however, if the information in the sign is stored as unstructured text in the map database. It is important that the Back Seat Driver understand what the sign says, not simply utter the words. There are two reasons for this. First, the Back Seat Driver's internal representation for text is preferably based on syntactic structure, not text strings. Second, the objects mentioned in the signs (cities and roads) should be entered into the discourse model to become salient for future reference. The Back Seat Driver should parse sign text by separating it into tokens delimited by commas and the word "and", and then attempt to recognize objects on the map (street names, cities, neighborhoods) from these tokens. When recognition fails, the token cannot be entered into the discourse model. When parsing fails, the spoken output will have incorrect grammar.

The Back Seat Driver does not assume that the driver will recognize the place to act (e.g. by seeing a street sign) so the driver must be told when (or where) to act. The Back Seat Driver uses timing ("Take a left here") when the driver has reached the place to act. The working prototypes calculate the place to speak by finding a distance from the intersection which is $v * (t_{speak} + t_{reaction})$, where t_{speak} is the time to speak the utterance and $t_{reaction}$ is the driver's reaction time. The time to speak depends on the number of words in the utterance. (The Dectalk synthesizer used in the prototypes speaks 180 words per minute.) Reaction time can be estimated to be two seconds.

The Back Seat driver speaks autonomously, but preferably provides means to allow it to speak on demand. The driver at any time should be able to ask for instructions immediately by, for example, pushing buttons, representing "What next?" and "What now?". In addition, while following a route, a driver should be able to ask to hear the total length of the route and the remaining distance. A driver should also be able to ask to hear the name of the street the car is on and the compass direction the car is headed.

In order to generate more fluent text, the Back Seat Driver preferably keeps track of what has been mentioned. Some instructions are obvious after having been given. If the system tells the driver to go straight through a set of lights, there is no reason to repeat the instruction when actually at the lights. This is in contrast with a turn, where the driver hears advance instructions to know what to do, and immediate instructions to know when to do it. This can be important, for if the driver hears "go straight through the lights" twice, she may try to go through two sets of lights. To implement this, each instruction should be able to determine whether it is obvious after having been given once. When it is time to speak the instruction, if the instruction has already been given, and it is obvious once spoken, then it should not be spoken again.

The Back Seat Driver preferably retains a history of the route. This allows it to generate cue phrases for the instructions. If the route calls for doing the same thing twice in a row, the system uses the work "another" to indicate this doubling. This is important for polite behavior. If a passenger were to give a driver instructions by just saying "Take a right. Take a right. Take a left. Take a right.", pronouncing each the same, the passenger would be judged to be rude. The passenger's speech is not acknowledging the driver's actions or history. There are two ways for the passenger to acknowledge

5,177,685

19

the driver's work: using cue words ("Take a right. Take another right. Now take a left."), or using intonation. However, some speech synthesizers, such as the DecTalk used in the prototypes, does not support flexible control of intonation, so cue words are the only possibility.

The Back Seat Driver preferably is able to warn the driver about dangers which can be inferred from knowledge of the road network. These dangers include driving above the speed limit, driving the wrong way on a one-way street, driving too fast for an upcoming curve, driving on a one-way street that becomes two-way ahead, merging traffic, "blind" driveways ahead, speed traps, poorly repaired roads, potholes, and dangerous intersections. The Back Seat Driver preferably attempts to determine hazards by reasoning about road conditions rather than requiring them to be built in to the map database.

Lane advice includes telling the driver which lane to get into (or stay out of) when applicable. The system gives lane advice as part of the instruction when approaching an intersection where it matters. The instruction may also include advice about what lane to be in after the intersection, in preparation for the next act.

Speed advice includes warning the driver to slow down if she is travelling too fast to safely negotiate a turn. The limiting factor for angular acceleration is the driver, not the cornering ability of the car. Research has shown that the average driver will accept no more than 0.1 G radial acceleration. Radial acceleration is v^2/r where r is the turning radius of the turn. The Back Seat Driver knows the geometry of the road, so it can predict the maximum tolerable velocity for the turn. It need not tell the driver about this speed (the driver will choose a comfortable speed without being told), but it should estimate the distance required to decelerate, and tell the driver to slow down early enough to do this gently.

If the driver leaves the route, the Back Seat Driver immediately informs the driver and begins to plan a new route. Telling the driver what she did wrong prepares her for hearing new instructions, and perhaps helps her learn to better interpret the style of language that the Back Seat Driver uses.

To describe an error, the Back Seat Driver needs to look back to the last action that the driver failed to perform. It should utter a description of this action, saying e.g. "Oops, I meant for you to take a right," which does not blame the driver as in e.g. "You made a mistake. You should have taken a right." A driver might leave the route deliberately, or the error could be system's, not the drivers.

Errors will occur due to inaccuracies in the location system. The Back Seat Driver is preferably able to model the uncertainty of a position. For instance, when two roads diverge at a narrow angle, it will be unable to distinguish which was taken until some distance passes. It should probably assume that the driver made the correct choice rather than taking the risk of making a false accusation. If it reaches a place where the difference is crucial, yet unknown, it is probably better for it to confess its uncertainty, and say something like "I'm not quite sure where we are, but if you can take a right here, do it, and if not, keep going, and I'll figure things out better in a minute." Or it might ask the driver to pull over and stop (if the driver is at a place where that is safe) to allow for a better position estimate by averaging a few successive estimates.

20

Errors will also occur if the database is somewhat out of date. The Back Seat Driver can regain at least a little confidence by how it explains the mistake. Suppose that the Back Seat Driver intends the driver to turn onto "Apple" Street, and says "Take a right at the next light". Unbeknownst to it, a new traffic light has been installed at "Pear" Street, so the driver turns there. It is somewhat confusing if the Back Seat Driver says "I meant for you to go straight," because the driver may think that the program has not been consistent. A better message would be "I did not mean for you to turn onto Pear. I thought that the next set of lights was at Apple Street."

While the driver is following a route, the Back Seat Driver preferably adopts a persistent goal of keeping the user reassured about her progress and the system's reliability. If the Back Seat Driver were a human, this might be unnecessary, since the driver could see for herself whether the navigator was awake and attending to the road and driver. But the driver cannot see the Back Seat Driver and so needs some periodic evidence that the system is still there. One piece of evidence is the safety warnings the system gives. But if all is going well, there will not be any. Other kinds of evidence that things are going well should be provided. When the user completes an action, the Back Seat Driver can acknowledge the driver's correct action, saying something like "nice work" or "good". Also, insignificant remarks about the roads nearby, the weather and so on, can be provided. The driver then assumes that everything is going well, for otherwise the Back Seat Driver would not speak of trivial matters.

The Back Seat Driver should know about the knowledge and desires of its driver, and act differently because of this knowledge. This knowledge is preferably incorporated in a user-model.

For driver properties which do not change or change very slowly, such as colorblindness, or visual or aural acuity, it is acceptable for the Back Seat Driver to ask the user for such knowledge. However, for other driver properties, the Back Seat Driver preferably acquires a model of the user automatically, without asking or having to be told. For example, the Back Seat Driver could learn the driver's reaction time by measuring the time between its speech and the driver's operation of the controls.

The Back Seat Driver preferably learns the style of instruction giving appropriate for the driver. To learn the driver's preferences for route description requires either observation of the driver herself giving instructions or getting feedback from the driver about the instructions the system provides.

The driver can provide feedback about the suitability of the Back Seat Driver's instructions either explicitly or implicitly. One explicit indication of comprehension is how often the driver hits the "what now?" button. The system might collect long term statistics on the types of intersections where the user most often requests help, and begin to offer instructions without being asked. Just as the user can ask for more talking with the "what now" button, the Back Seat Driver should provide a "shut up" button (or other means to the same effect). The Back Seat Driver could also learn the effectiveness of its directions simply by watching the driver's performance—in particular, her errors. In this way, it can learn which cues are most effective.

Another opportunity for learning the driver's style is during the acquisition of speech recognition templates

(for user-dependent speech recognition for driver input means, described below). The new user should play the role of a "back seat driver" and give instructions, while in a car, for some route. The instructions must be given while driving either a real car or a close simulation because the form of static driving instructions is much different from real time instructions. Given some a priori knowledge about the ways that a route can be described, it is not impossible that the system could understand the instructions, and infer style from it. A difficulty here is that if the driver knows the route well, many things will seem obvious to her that would not be obvious to another person.

If the Back Seat Driver knows what the driver knows about the city, it can give better directions. A user who knows about a city no longer need instructions, she needs information about structure. The object description system preferably provides novice users a process description which emphasizes casual connections, and experts structural descriptions. Experts do not need the casual information, they can derive it for themselves.

The attributes of the user-model preferably include: route-preference—does the driver want the fastest, shortest, or simplest route?

reassurance-period—how often should the program speak to the driver?

use-names—should the program tell the driver the names of passing streets?

congratulate-after-act—should the program make an explicit acknowledgment of correctness to the driver after each act?

obvious-to-cross-major—can the program assume that the driver will continue straight across a major intersection without being told explicitly to do so?

scofflaw—does the driver want to be warned when she is speeding?

daredevil—does the driver want warnings when driving dangerously fast?

distance-lowpass—does the driver want to be told the distance to the next action (in yards or miles, as appropriate). Most drivers do not understand distances in tenths of miles, so by default the program mentions only those distances that exceed one half mile.

The functions of the user-model preferably include: obvious-next-segment—given a current position, is there a unique segment such that it is almost certain the driver will go there, without being told to do so? at-major-intersection—is the current intersection one that the driver would call "major"?

extrapolate-path—try to predict the path the driver will follow in the next N seconds, assuming she does only what is obvious.

maximum-safe-speed—calculate the maximum speed at which the driver can get through an intersection. This calculation is based on finding the segment with the greatest radius of turn, and then calculating the largest speed the vehicle could have while making that turn without undergoing unacceptable sideways acceleration.

For the Back Seat Driver to decide what to say and when to say it, it preferably has a model of the vehicle performance. It must know, for example, how slowly the car should be going in order to safely make a turn. A suitably instrumented car could also measure the coefficient of friction by comparing the applied braking force and the resulting deceleration. This would allow it to adjust the time factors used in deciding when to speak.

The Back Seat Driver should understand the driver's plans and goals. When a driver enters a destination address, she is telling the system that she has the goal of getting to that address. The Back Seat Driver might guess at higher level plans from knowledge about the destination, and take actions to assist the driver with those plans. To do this, it must know what kind of thing is at the destination address. For instance, if the address provided is that of a store, the Back Seat Driver can guess that the driver is going there to purchase something, or at least to do business there. It could check the hours that the store is open.

The Back Seat Driver should help drivers to understand the route it gives. This would make the system more pleasant to use. It would also make it easier to follow routes because a driver who understands the route and the city will use that knowledge to help interpret the commands Back Seat Driver gives. A route should fit into a larger model of the city. This means that the Back Seat Driver itself must have a model of the city and should speak of the route in terms that relate it to the city. There are several opportunities to do this. At the beginning of the route, the driver might hear an overview of the route, naming the major paths followed and neighborhoods crossed. During the route, locations could be described not just as street address but in larger units of neighborhoods and districts. Orienting information can be included in instructions, or it might come between instructions, as a passing comment.

There are some additional services that the Back Seat Driver could easily provide. It should be able to give the location of a place without giving directions, it should be able to give the directions all at once, and it should be able to give directions between any two places. A driver might want to use these to tell someone else how to get to where they are.

The Back Seat Driver should be able to communicate with the outside world if the outside world is equipped to talk to it. For instance, after determining that a given parking garage is the closest or most convenient to the current destination, the Back Seat Driver could automatically phone or radio the garage and reserve a space.

The Back Seat Driver should be running on a computer embedded in the car, so that it can get more and better information about the state of the car and driver. For instance, when the next instruction is a turn, the Back Seat Driver should notice whether and when the driver turns on the turn signals. If the driver applies them too soon, it is possible (but not certain) that the driver has underestimated the distance to the turn; if applied at the "right time" then the system can take that the action has been understood; if never applied, then the driver has either misunderstood, or is driving hazardously.

The Back Seat Driver should also be integrated into the car's audio system, rather than having separate systems for voice and music. Furthermore, it should pay attention to what the driver is listening to. If the driver is listening to the radio, or playing a CD (or using a cellular telephone) the program should try to speak less often, on the grounds that the driver has implicitly indicated a preference for what to listen to. The program should suppress reminders and historical notes altogether. When it must speak, it should borrow the audio channel rather than trying to speak over it. The Back Seat Driver should also be aware of the driver's use of other controls in the car. It should defer speech

while the driver is adjusting the heat or the mirrors, for example, and suppress speaking altogether if the car makes sudden extreme changes in velocity. A driver trying to cope with an emergency situation does not need another distraction.

The discourse model preferred for the Back Seat Driver is a partial implementation of the discourse theory described by B. J. Grosz and C. L. Sidner ("Attention, intentions, and the structure of discourse" in *Computational Linguistics*, 12(3):175-204, 1986) and the theory of intonational meaning described by J. Hirschberg and J. Pierrehumbert ("The intonational structuring of discourse" in *Proceedings of the Association for Computational Linguistics*, 136-144, July 1986). Both of these articles are herein incorporated by reference. This model allows the program (or programmer) to create and manipulate discourse segments. The program specifies the contents of the discourse segment (both the syntactic form and the list of objects referenced) and the implementation of the model does the following: generates prosodic features to convey discourse structure; inserts discourse segment into intentional structure; and maintains attentional structure—adding new objects when mentioned and removing old objects when replaced. In addition it includes some useful low-level tools for natural language generation: search of attentional structure for occurrence of co-referential objects; conjugation of verbs; generation of contracted forms; and, combination of sentences as "justification" sentences (e.g. "get in the right lane because you are going to take a right.") and sequential sentences ("Go three blocks, then turn left"). In order to use the discourse package the programmer must explicitly declare all semantic types used by the program, so in this case there are declarations for all objects which pertain to driving, such as street names, bridges, rotaries, stop lights and so on.

SPEECH GENERATOR

In the working prototypes of the Back Seat Driver, speech generation is performed by Dectalk, a commercial text-to-speech speech synthesizer, which is cabled to the main computing apparatus.

An alternative to synthesized speech is digitized speech, which is easier to understand than synthetic speech. Digitized speech, however, requires a great deal of storage space. There are more than 2000 different street names in Boston. Allowing another 500 words for the actual instructions, and assuming an average duration of 0.5 seconds for each word, coding this vocabulary at 64 kilobits per second would require 10 megabytes of speech storage. Given a Back Seat Driver that uses a CD-ROM for the map, the digitized speech could be stored on the disk as well. Coded speech would be more intelligible than synthesized speech, and less costly, but not as flexible. For, example, it would be impossible to read electronic mail using only stored vocabulary speech.

The generated speech is spoken to the driver through some kind of speaker system in the car. In a simple embodiment, the speaker system of the car radio is used.

DRIVER INPUT MEANS

Means for the driver to communicate with the back-seat driver are required. For example, the driver must be able to enter destination addresses, request instructions or a repeat of instruction, and inform the Back Seat driver when an instruction cannot be carried out

for some reason. In embodiments where the computing apparatus is installed in the automobile, a computer keyboard can be adapted to provide this communication means.

- 5 In one working prototype of the Back Seat Driver, the computing apparatus is not installed in the automobile, but is accessed through a cellular telephone. In this embodiment, the driver communicates with the Back Seat Driver by using the cellular telephone keypad.
- 10 Address entry is achieved by first entering the digits, then a number sign, then spelling the street name using the letters on the telephone keypad. The letters "Q" and "Z" are on the "6" and "9" keys, respectively, and the space character is on "1", which is otherwise unused.
- 15 These keys are sufficient to spell any street name in Boston. The spelling rules can be easily expanded to enter street names with other characters in them, for example, "4th Street".

In the implementation, spelling a street name requires only one button push for each letter, even though there are three letters on each key. This is because of the redundancy in street names, which are pronounceable words, not arbitrary strings. There are 37 pairs of street names in Boston with the same "spelling" in the reduced "alphabet". An example is "Flint" and "Eliot", both encoded as "35468". This is only one percent of the 2628 names of streets in Boston, so collisions are rare. This technique appears workable even for larger sets of names. When the entire word list of the Brown corpus is encoded, there are still only 1095 collisions in the 19,837 words (5.5%).

If a name collision occurs, the Back Seat Driver reads the list of possibilities, and asks the driver which one was meant. This is very rare. A more common problem is that street names are duplicated. When this happens, the Back Seat Driver asks the user a series of questions to reduce the list to a single choice. It tries to ask the fewest questions possible. It asks the user to choose from a list of street types, if that is sufficient to resolve the question, and otherwise from a list of the containing cities (or neighborhoods, if there are two instances within a single city). To select from a list, the Back Seat Driver reads the contents, asking the user to push a button when the desired choice is read.

The Back Seat Driver would be much easier to use if the driver could simply talk to it instead of using a keyboard or keypad. Speech recognition could be used for driver input means, however, address entry is a difficult task for speech recognition for the same reason it is hard for a human to understand machine speech—there are few constraints on a name. No speech recognizer available today can handle a 3000 word vocabulary with acceptable error rates. The car would also have to be stopped while the driver was speaking, because noise in moving cars for frequencies below 400 Hz can exceed 80 dB.

Back Seat Driver could also use speech recognition to replace the "What now?" and "What next?" buttons. This is a more tolerant application for speech recognition because there are fewer words to recognize.

SYSTEM PROCESSES

The Back Seat Driver carries out three separate tasks, each of which is executed by its own process. All processes share the same address space, so all variables and functions are accessible in every process, and no special mechanism for interprocedure call is required. Where necessary for synchronization, Back Seat Driver uses

5,177,685

25

queues or locks. All three processes are simple, infinite loops. The system processes are illustrated in FIG. 2.

The user process is the main process of the Back Seat Driver. It is this process which finds routes and gives instructions to the driver. The user process manages a list of goals. Each time around the loop, it selects a goal to work on, and does something to achieve the goal, if possible. The user process is connected to the speech generator, since that is how it talks to the driver.

The navigator process maintains an estimate of the current position and velocity of the car. It is connected to the position sensor by a serial line. Preferably, packets arrive from the position sensor several times a second. The navigator converts the data in the packets from the position sensor format to the format used by the Back Seat Driver.

There are two minor processes which assist the navigator process: The average speed process computes the running average speed of the vehicle over the last five seconds. It could be made part of the navigator process, but is distinct because it is more convenient that way. The position sensor monitor process keeps track of how often packets arrive. If packets do not arrive when scheduled, it should set a flag to indicate this to inform the driver if the position sensor ceases to work properly.

The control process is responsible for controlling the Back Seat Driver as a whole. The control process is connected to driver input means for entering, for example, the destination and requesting additional instructions while driving (e.g. the "What now?", "What next?" and "I can't do it" features.) Other functions of the control process are useful in a research prototype, but should not be required in a commercial embodiment of the Back Seat Driver. One such function is debug-

ging.

The user process is a goal-driven perpetual loop which seeks to use the resources available to it to satisfy as many goals as possible as quickly as possible, devoting resources first to those goals which are of greatest importance. There are two aspects to this process, goal structures (which names goals) and goal-functions (which tell how to accomplish them). A goal is just a name, a priority (a number), and a set of slots (parameters). Thus for instance a typical goal would be (GET-TO-PLACE<140 Elm Street>), where the goal has one slot, namely the destination. A goal-function is a function which is possibly able to achieve a goal. When a new type of goal is defined, the programmer also tells the system which goal functions can possibly meet it, and later, when the system tries to accomplish a goal it selects from this list.

The goal loop is a three step process. 1) Check to see whether there are any newly added goals. The driver can add a goal by hitting a key, and the system can also add goals. 2) Find the most important goal to work on. 3) Work on that goal. In general, systems should use resources in the most efficient manner possible. For the Back Seat Driver, the only resource is speaking time. The only way the Back Seat Driver can accomplish any of its goals is by speaking. Speech is a resource because the program can only say one thing at a time, and speaking takes a finite time. It is also important to note that spoken utterance has a useful effect only if completely spoken, so it is not helpful to begin to speak if there is not time to complete the speech.

The protocol for a goal function preferably includes the following:

26

progressable?—Is the goal able to accomplish anything at this time?

resource-used—If it runs now, what resources will it want to use?

maximum-time-of-resource—If it runs now, how long (in seconds) will it need each resource?

minimum-time-to-resource—The minimum time that it can estimate until it may again need this resource, and the priority of its use at that time.

In the working prototypes of the Back Seat Driver, the list of all goals is stored in the global variable *goals*. The goal loop and goal structures are defined in the file goals.lisp. The various goals and goal functions of the Back Seat Driver are defined in the files main.lisp, route-goals.lisp, and get-to-place.lisp. All goals which use speech are built from the speech-goal object defined in speech-goal.lisp. The speech resource itself is defined in speech-resource.lisp.

The goal or function which gets a user to a destination is called GET-TO-PLACE. An explanation of this goal will illustrate the goal mechanism in more detail, as well as illustrate how this most important function of Back Seat Driver is implemented. The goal GET-TO-PLACE, has two slots, destination which is the location the user wants to get to, and route which is the route the Back Seat Driver intends to use to get there.

The driver adds the goal to the system goal list by striking a key. When the goal is first created, the destination is not known (the destination slot is empty), so the goal function for GET-TO-PLACE creates a sub-goal, GET-DESTINATION, and adds it to the goal list. Now there are two goals on the goal list, GET-TO-PLACE and GET-DESTINATION, but only the second is progressable, because when a goal has a sub-goal, it is not allowed to run until the sub-goal finishes. Therefore, the only progressable goal is GET-DESTINATION, which attempts to get a destination by asking the user to enter an address. If the user fails to do so, the subgoal fails, which in turn causes GET-TO-PLACE to fail, and the Back Seat Driver says "Travel cancelled". Otherwise, it writes the destination into the destination slot of the GET-TO-PLACE goal. Now that the sub-goal is complete, GET-TO-PLACE can once again make progress. This time it finds that the route slot is empty, and again calls for the sub-goal GET-ROUTE, which calculates a route. When this is complete a third subgoal is invoked, namely FOLLOW-ROUTE.

The goal function for FOLLOW-ROUTE gets the driver to the destination by speaking instructions. If something goes wrong (for example if the driver makes a mistake) then the subgoal fails. But this does not make GET-TO-PLACE give up. Instead, it erases the route slot, and simply finds a new route, and then tries FOLLOW-ROUTE again. This continues, no matter how many times things go astray, until either FOLLOW-ROUTE succeeds, or the driver cancels the trip.

The goal FIND-SERVICE is similar to GET-TO-PLACE except the driver selects a kind of service (for example, a gas station), and then the Back Seat Driver finds the closest provider of that service, and then finds a route to it. Following that route is done by FOLLOW-ROUTE in the same way as for GET-TO-PLACE.

The FOLLOW-ROUTE goal function gets the user to her destination by giving spoken instructions. There are several reasons it might speak:
at the beginning, to alert the driver

5,177,685

27

- to give an instruction in advance, so the driver will be ready
- to give an instruction when it is time to do it
- to confirm that the driver has correctly carried out an instruction
- to inform the driver of her arrival at the destination
- to reassure the driver that she is still on route
- to inform the driver of a mistake
- to warn the driver that she is driving so fast that the program cannot keep up.

FOLLOW-ROUTE decides the next reason for speaking by first trying to locate the current position on the path. If the position is not on the path (more precisely, if the current segment does not occur anywhere on the path) then the driver has left the path (or the position sensor has made an error). Otherwise, FOLLOW-ROUTE determines what instruction must be next executed by calling the function next-driver-instruction.

The goal function protocol requires that FOLLOW-ROUTE support the goal function minimum-time-to-resource, which estimates the minimum time until FOLLOW-ROUTE will next speak. This time depends upon the reason for the next speaking. FOLLOW-ROUTE speaks immediately when beginning, confirming, warning, or finishing the route. When the driver is off the route, FOLLOW-ROUTE waits a few seconds before speaking, just in case the driver's departure from the route is actually a temporary error by the position sensor.

Given that the driver is on the path, FOLLOW-ROUTE determines when to speak by calculating the position where it must begin speaking the instruction text, then estimating the time required to reach that position at the driver's current speed. As the driver's speed changes, so will this estimated time. The position to begin speaking is calculated by first finding the position where the instruction is executed, then moving back a distance to allow the Back Seat Driver time to speak the text and the driver to react to it.

The Back Seat Driver can also give instructions in advance, if desired. It does this in much the same way, except that it adds an additional number of seconds (normally thirty) to the time estimate, and so begins to speak much sooner. When it gives instructions in advance the instruction text is longer because the program has more time to speak.

When the driver leaves the route FOLLOW-ROUTE starts a timer. If the driver has not returned to the route by the time the timer goes off (at present, two seconds) then FOLLOW-ROUTE checks for a possible mistake. In describing the mistake, it attempts to characterize what the driver actually did as well as what the program intended the driver to do. It is able to do this because in the main loop it stored the last position that the driver was on when last on the route.

Goals may interrupt lower priority goals by requesting the speech resource to interrupt the lower priority goal. Interruption stops the speech-synthesizer immediately. The interrupted goal is informed of the interruption, and can react as it chooses. There is no way for the goal to know whether any of its words were actually spoken, so it has to start all over. Most goals attempt to run again as soon as possible. The assumption is that the interruption occurred because the user started some higher priority goal after learning how to do so through the help command.

28

The system treats "repeat the last statement" as a goal, rather than as a special purpose function, except that the importance of this goal is set to the value of the last goal spoken (the goal whose utterance is being repeated). This guarantees that if some more important goal desires to speak, it will be able to. A repetition of an utterance is no more important than it was originally.

Goals can be temporary or persistent. Temporary goals can be satisfied, but persistent goals never can be.

10 All system initiated goals are persistent. The system goals include warning the driver of dangers ahead (WARN-DRIVER) and informing the user of new electronic mail or other messages (if the computer apparatus of the Back Seat Driver is connected to the outside world). These goals can never be satisfied. The driver's safety should always be preserved and mail or messages can arrive at any time.

CELLULAR PHONE EMBODIMENT

20 The Back Seat Driver is preferably an in-car navigation system, but in some embodiments, it may be desirable to not have the entire computing apparatus installed in the car. This is the case if the computing apparatus is too large or if a number of cars are to share a 25 single computing apparatus.

For such embodiments, two cellular phones installed in the car can be used to transmit data from the car to the computing apparatus, and to receive voice from the speech generator in the computing apparatus. In this 30 embodiment, data from the position sensor installed in the automobile is sent through a cellular phone in the car equipped with a modem to a phone connected to the computing apparatus via a modem. The voice generating apparatus of the computing apparatus sends speech 35 over another phone to a second cellular phone installed in the automobile.

This embodiment has been implemented in a working prototype, using a large workstation computer (a Symbolics Lisp Machine). In this implementation, a position 40 sensor installed in the car estimates vehicle position, heading, and velocity, and sends a data packet, once per second, through a modem to the workstation. The workstation sends characters to a Dectalk speech synthesizer, which in turn sends voice over a second phone to the driver.

Nearly everyone who has used a cellular phone knows how noisy they are. Cross talk is common and noise bursts and signal loss make it hard to hear. A sufficiently bad noise burst can even cause the cellular 50 system to terminate the call. The problems for data transmission are even worse. By its very nature, cellular radio transmission is subject to multi-path interference, which causes periodic fades as the antenna moves in and out of anti-nodes. In addition to this fading, there is a 55 complete loss of audio signal for as long as 0.9 seconds when the phone switches from one cell site to another (hand off).

An attempt to use an ordinary (land-line) modem from the car was unsuccessful. In the prototype, a 60 Worldlink 1200 from Touchbase Systems was used in the car, with a Morrison and Dempsey AB1 data adapter, and an NEC P9100 phone, boosted to 3 watts. At the base station, both a Practical Peripherals 2400 and a Hayes Smartmodem 1200 were used. Even at 300 baud the connection was too noisy to use. Worse, connections seldom lasted more than five minutes. In all cases, the "loss of carrier" register (S10) was set to its maximum value, 20 seconds. Loss of carrier signal alone

5,177,685

29

cannot explain why the connections dropped. The modems were capable of tolerating a complete loss of audio for up to twenty seconds.

Better results were found using an error correcting modem (The "Bridge") made by the Spectrum Cellular Corporation. This modem uses a proprietary protocol (SPCL) for error correction. The Spectrum product virtually eliminated noise, at the price of a lower data transmission rate. According to the protocol, the transmitting modem groups characters into packets that include error correction bytes. If only few errors occur, the receiving modem repairs the data and acknowledge receipt. If there are many errors, the packet is retransmitted. If the sending modem has to retransmit too often it makes the packets smaller, on the assumption that a smaller packet has a better chance of success. This is less efficient, since packets have a fixed overhead, the percent of the channel used by data decreases. When conditions improve the modem increases packet size again. In theory, the modem can transmit at 120 characters per second, but tests made by recording the time required to receive the three characters of an odometer sequence demonstrated that the average value is closer to 30 characters per second. This sequence is transmitted once per second. The mean for durations for the three character sequences is 94 milliseconds, which is 31 milliseconds per character, or 32 characters per second.

Even with the cellular modem, calls are sometimes dropped. The call durations are usually long enough for a successful trip with the Back Seat Driver. Voice calls are dropped at about the same rate as data calls.

The protocol used by the Spectrum modem acknowledges all data transmitted. If the acknowledgment is not received, it retransmits the data until acknowledged. Under adverse conditions, this can result in an arbitrarily long delay. This is a problem when real-time data is transmitted. Observation of the Back Seat Driver shows that sometimes the system will "freeze" from one to ten seconds. During this time, the car of course continues to move. If these freezes occur near decision points, the driver may go past the intersection without being told what to do. At 20 miles per hour a car travels nearly 45 meters in five seconds. The navigation system in the car sends a packet once every second. Most packets arrive within a second, but a few are delayed, some by up to ten seconds. (These delays may also arise from delays at the workstation. Lisp Machines are not noted for real-time response.)

It would be better to have a protocol which guarantees to deliver data intact and free of errors, if it delivers it at all, but does not guarantee to deliver the data. Real time data is only valuable in real time, and time spent retransmitting old data is taken away from ever, more valuable data. Such a protocol modification is feasible for the Spectrum product.

What is claimed is:

1. An automobile navigation system which produces spoken instructions to direct a driver of an automobile to a destination in real time comprising:
computing apparatus for running and coordinating system processes.
driver input means functionally connected to said computing apparatus for entering data into said computing apparatus, said data including a desired destination.

30

a map database functionally connected to said computing apparatus which distinguishes between physical and legal connectivity.

position sensing apparatus installed in the automobile and functionally connected to said computing apparatus for providing said computing apparatus data for determining the automobile's current position,

a location system functionally connected to said computing apparatus for accepting data from said position sensing apparatus, for consulting said map database, and for determining the automobile's current position relative to the map database,
a route-finder functionally connected to said computing apparatus, for accepting the desired destination from said driver input means and the current position from said location system, for consulting said map database, and for computing a route to the destination,

a discourse generator functionally connected to said computing apparatus for accepting the current position from said location system and the route from said route finder, for consulting said map database, and for composing discourse including instructions and other messages for directing the driver to the destination from the current position.

a speech generator functionally connected to said discourse generator for generating speech from said discourse provided by said discourse generator, and

voice apparatus functionally connected to said speech generator for communicating said speech provided by said speech generator to said driver.

2. The automobile navigation system of claim 1 wherein said map database comprises a set of straight line segments and a set of nodes, each endpoint of each segment being a pointer to a node representing the coordinates of the endpoint and the set of other segments which are physically and legally connected to that endpoint.

3. The automobile navigation system of claim 1 wherein said map database is based on DIME files of the United States Census extended to represent physical and legal connectivity.

4. The automobile navigation system of claim 3 wherein said DIME file is further extended to distinguish bridges, underpasses, tunnels, rotaries, and access ramps from other street types.

5. The automobile navigation system of claim 1 wherein said map database is based on TIGER files of the United States Census and United States Geological Survey extended to represent physical and legal connectivity.

6. The automobile navigation system of claim 5 wherein said TIGER file is further extended to distinguish bridges, underpasses, tunnels, rotaries, and access ramps, from other street types.

7. The automobile navigation system of claim 1 wherein said map database comprises a three-dimensional representation of street topology.

8. The automobile navigation system of claim 1 wherein said map database includes measures of street quality.

9. The automobile navigation system of claim 1 wherein said map database distinguishes divided streets.

10. The automobile navigation system of claim 1 wherein said map database includes landmarks such as signs, traffic lights, stop signs and buildings.

5,177,685

31

11. The automobile navigation system of claim 1 wherein said map database includes lane information.
12. The automobile navigation system of claim 1 wherein said map database includes speed limits.
13. The automobile navigation system of claim 1 wherein said map database includes expected rate of travel.
14. The automobile navigation system of claim 1 wherein said map database includes time-dependent legal connectivity.
15. The automobile navigation system of claim 1 wherein said map database includes turn difficulty.
16. The automobile navigation system of claim 1 wherein said map database includes vehicle street, lane, and height restrictions.
17. The automobile navigation system of claim 1 wherein said map database includes traffic light cycles.
18. The automobile navigation system of claim 1 wherein said map database distinguishes where right turn on red is allowed.
19. The automobile navigation system of claim 1 wherein said map database includes a database of service locations.
20. The automobile navigation system of claim 1 wherein said map database includes a listing of famous places by name.
21. The automobile navigation system of claim 1 further comprising means for updating said map database.
22. The automobile navigation system of claim 1 further comprising means for updating said map database by radio broadcast.
23. The automobile navigation system of claim 1 wherein the map has minimum accuracy of 10 meters.
24. The automobile navigation system of claim 1 wherein said route finder is based on a best-first search algorithm.
25. The automobile navigation system of claim 1 wherein said route finder is based on an A* algorithm.
26. The automobile navigation system of claim 1 wherein said route finder is based on an A* algorithm modified to find a route in less time.
27. The automobile navigation system of claim 1 wherein said route finder is adapted to find a best route according to any one of three cost metrics: distance, speed, simplicity.
28. The automobile navigation system of claim 1 wherein said route finder is adapted to calculate a new route if the driver or vehicle navigation system makes an error or if the route is unnavigable due to unforeseen circumstances, wherein said new route does not simply backtrack to the point of the error if a better route from the current location exists.
29. The automobile navigation system of claim 1 wherein said route finder is adapted to calculate a new route while the automobile is in motion, wherein said new route will begin from the location of the automobile at the time the calculation of the new route is completed.
30. The automobile navigation system of claim 29 wherein an estimated time to find a new route is multiplied by the velocity of the automobile to calculate the position from which the new route should start.
31. The automobile navigation system of claim 30 wherein said estimated time to find a new route is calculated by multiplying the distance between the starting and ending points of the new route by a constant.

32

32. The automobile navigation system of claim 1 wherein said location system is a position-keeping (dead-reckoning) system.
33. The automobile navigation system of claim 1 wherein said location system is a hybrid of position-keeping and position-finding systems.
34. The automobile navigation system of claim 1 wherein said location system employs map matching.
35. The automobile navigation system of claim 1 wherein said position sensing apparatus comprises displacement and direction sensors installed in the automobile.
36. The automobile navigation system of claim 1 wherein said position sensing apparatus measures displacement with an odometer.
37. The automobile navigation system of claim 1 wherein said position sensing apparatus measures direction with a magnetic compass.
38. The automobile navigation system of claim 1 wherein said position sensing apparatus measures direction by monitoring the turning of the steering wheel.
39. The automobile navigation system of claim 1 wherein said position sensing apparatus measures direction with a differential odometer.
40. The automobile navigation system of claim 1 wherein said position sensing apparatus measures direction with a gyroscope.
41. The automobile navigation system of claim 1 wherein said discourse generator is based on an object-oriented programming methodology.
42. The automobile navigation system of claim 1 wherein each intersection in a route is classified into one type in a taxonomy of intersection types, and the disclosure generated in relation to each said intersection depends on its type.
43. The automobile navigation system of claim 42 wherein said taxonomy of intersection types includes continue, forced-turn, U-turn, enter, exit, onto-rotary, stay-on-rotary, exit-rotary, fork, turn, and stop.
44. The automobile navigation system of claim 42 wherein said discourse generated further depends on a description function for each intersection type which generates a description given the length and tense of the desired description and the position along the route from which an instruction is to be given.
45. The automobile navigation system of claim 1 wherein said discourse generated comprises a long description of an act given substantially before the act is to be performed and a short description given at the time the act is to be performed.
46. The automobile navigation system of claim 45 wherein said long descriptions includes cues.
47. The automobile navigation system of claim 46 wherein said cue is a landmark.
48. The automobile navigation system of claim 1 wherein said driver input means includes means for said driver to demand immediate instructions, or clarification or repetition of instructions already provided.
49. The automobile navigation system of claim 1 wherein said driver input means includes means for said driver to indicate to said automobile navigation system that a given instruction provided by said system is impossible to complete for some reason and that a new route must be calculated.
50. The automobile navigation system of claim 1 wherein said driver input means comprises a voice recognition system to allow at least some driver input to be spoken.

5,177,685

33

51. The automobile navigation system of claim 1 wherein said automobile navigation system records a history of the route and the discourse already generated and uses this knowledge to generate cues for future discourse and make future discourse more understandable.

52. The automobile navigation system of claim 1 wherein said automobile navigation system warns drivers of dangers inferred from knowledge of the road network.

53. The automobile navigation system of claim 1 wherein said automobile navigation system informs a driver if an error has been made as detected by the location system.

54. The automobile navigation system of claim 15 wherein said discourse generator is responsive to a user-model stored in said computing apparatus to customize discourse to the requirements and preferences of said driver.

34

55. The automobile navigation system of claim 1 wherein said speech generator is a speech synthesizer.

56. The automobile navigation system of claim 1 wherein said speech generator uses digitized speech.

57. The automobile navigation system of claim 1 wherein said computing apparatus is not installed in the automobile, and said automobile navigation system further comprises means for communication between said computing apparatus and the automobile navigation system components installed in the automobile.

58. The automobile navigation system of claim 57 wherein said means for communication is two cellular phones in said automobile, one of which is connected to a modem, and two phones connected to said computing apparatus, one of which is connected to a modem, whereby one data channel and one voice channel between said automobile and said computing apparatus is created.

* * * * *

20

25

30

35

40

45

50

55

60

65

EXHIBIT

2

IW 1415424

THE UNITED STATES OF AMERICA

TO ALL TO WHOM THESE PRESENTS SHALL COME:

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office

February 07, 2006

THIS IS TO CERTIFY THAT ANNEXED IS A TRUE COPY FROM THE
RECORDS OF THIS OFFICE OF THE FILE WRAPPER AND CONTENTS
OF:

APPLICATION NUMBER: 07/565,274

FILING DATE: August 09, 1990

PATENT NUMBER: 5,177,685

ISSUE DATE: January 05, 1993

By Authority of the
Under Secretary of Commerce for Intellectual Property
and Director of the United States Patent and Trademark Office



A handwritten signature in cursive script, appearing to read "N. Williams".
N. WILLIAMS
Certifying Officer

PART (2) OF (3) PART(S)
0001

10 093274		PATENT DATE JAN 05 1993		PATENT NUMBER 5177685	
ITEM NUMBER 77563274	FILED DATE 1/9/90	CLASS 364	SUBCLASS	GROUP 2	UNIT EXAMINER SCHMANDT
NAME OF INVENTOR JAMES R. DAVIS, NORTH CAMBRIDGE, MA; CHRISTOPHER M. SCHMANDT, MILTON, MA.					
CONTINUING DATA*** VERIFIED <i>NONE</i>					
FOREIGN/PCT APPLICATIONS*** VERIFIED <i>NONE</i>					
CERTIFICATE MAILED 1-15-96 OF CORRECT					
FOREIGN FILING LICENSE GRANTED 6/9/22/90					
Right(s) claimed: USC 113 conditions met: Filed and Acknowledged	<input type="checkbox"/> yes <input checked="" type="checkbox"/> no <input type="checkbox"/> yes <input checked="" type="checkbox"/> no Examiner's Initials	AS FILED →	STATE OR COUNTRY MA	SHEETS DRWGS 2	TOTAL CLAIMS 58
		INDEP CLAIMS 1	FILING FEE RECEIVED \$ 946.00	ATTORNEY DOCKET NO.	
SAM PASTERNACK CHOATE, HALL & STEWART EXCHANGE PLACE 100 STATE STREET BOSTON, MA 02109					
AUTOMOBILE NAVIGATION SYSTEM Using Real Time Spoken Instructions					
U.S. DEPT. OF COMMERCE - PATENT & TRADEMARK OFFICE - PTAB-901 rev 7/07					
PARTS OF APPLICATION FILED SEPARATELY					
NON-US CITATION AND MAILED		ISSUE DATE PREPARED FOR ISSUE		CLAIMS ALLOWED	
10/20/97 (10-20-97)		10/20/97 (10-20-97)		TOTAL CLAIMS 58	
Assistant Examiner Docket Clerk				DRAWINGS	
ISSUE FEES (10-20-97)				Sheets Drawings	
10/20/97 (10-20-97)				Fig. Drawings	
10/20/97 (10-20-97)				Photocopies	
10/20/97 (10-20-97)				Fees	
10/20/97 (10-20-97)				Other	
10/20/97 (10-20-97)				TOTAL	
10/20/97 (10-20-97)				58	
ISSUE CLASSIFICATION					
C10B 43/00		Subclass 443		Batch Number 520	
WARNING: The information disclosed herein may be restricted. Unauthorized disclosure may be prohibited by the United States Code Title 35, Sections 122, 161 and 361. Possession outside the U.S. Patent & Trademark Office is restricted to authorized employees and contractors only.					
DP02					

demand. The Back Seat Driver has the ability to construct a new utterance with the same intent, but not necessarily the same words, as a previous message. Synthetic speech being sometimes hard to understand, the Back Seat Driver chooses its words to provide redundancy in its utterances.

An actual working prototype of the Back Seat Driver has been implemented. It has successfully guided drivers unfamiliar with Cambridge, Massachusetts to their destinations. It is easy to foresee a practical implementation in the future.

D&C Brief Description of the Drawing

Fig. 1 illustrates schematically the basic functional components of the Back Seat Driver in its preferred embodiment.

Fig. 2 illustrates the system processes of the preferred embodiment of the Back Seat Driver.

D&C Description of the Preferred Embodiment

The automobile navigation system according to the present invention is illustrated schematically in Fig. 1. The heart of the system is a computing apparatus 10 comprising a vehicle location system 12, a map database 14, a route finder 16, a discourse generator 18, and a speech generator 20. Driver input means 22 allows the driver to input to the computing apparatus 10 information such as a desired destination. A position sensor 24 measures automobile movement (magnitude and direction) and sends data to the location system 12 which tracks the position of the automobile on the map. The route finder 16 calculates a route to the destination. Based on the current position of the automobile and the route, the discourse generator 18 composes driving instructions and other messages according to a discourse model in real time as they are needed. The instructions and messages are sent to the speech generator 20 which conveys them to the driver by means of a speaker system 26. The speaker system may be that of the car's radio.

In Fig. 1, the computing apparatus is illustrated as a single entity. However, in other



UNITED STATES DEPARTMENT OF COMMERCE
Patent and Trademark Office
Address: COMMISSIONER OF PATENTS AND TRADEMARKS
Washington, D.C. 20231

SERIAL NUMBER	FILING DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO.
07/565,274	08/09/90	DAVIS	J

SAM PASTERNACK
CHOATE, HALL & STEWART
EXCHANGE PLACE
53 STATE STREET
BOSTON, MA 02109

000

DATE MAILED: 08/28/90

**NOTICE TO FILE MISSING PARTS OF APPLICATION—
FILING DATE GRANTED**

A filing date has been granted to this application. However, the following parts are missing.

If all missing parts are filed within the period set below, the total amount owed by applicant as a large entity, small entity (verified statement filed), is \$ ~~104~~.

1. The statutory basic filing fee is missing. insufficient. Applicant as a large entity, small entity, must submit ~~\$ 104~~ to complete the basic filing fee and MUST ALSO SUBMIT THE SURCHARGE AS INDICATED BELOW.

2. Additional claim fees of \$ ~~200.00~~ as a large entity, small entity, including any required multiple dependent claim fee, are required. Applicant must submit the additional claim fees or cancel the additional claims for which fees are due. NO SURCHARGE IS REQUIRED FOR THIS ITEM.

3. The oath or declaration:
 is missing.
 does not cover items omitted at the time of execution.

An oath or declaration in compliance with 37 CFR 1.63, identifying the application by the above Serial Number and Filing Date is required. A SURCHARGE MUST ALSO BE SUBMITTED AS INDICATED BELOW.

4. The oath or declaration does not identify the application to which it applies. An oath or declaration in compliance with 37 CFR 1.63 identifying the application by the above Serial Number and Filing Date is required. A SURCHARGE MUST ALSO BE SUBMITTED AS INDICATED BELOW.

5. The signature to the oath or declaration is: missing; a reproduction; by a person other than the inventor or a person qualified under 37 CFR 1.42, 1.43, or 1.47. A properly signed oath or declaration in compliance with 37 CFR 1.63, identifying the application by the above Serial Number and Filing Date is required. A SURCHARGE MUST ALSO BE SUBMITTED AS INDICATED BELOW.

6. The signature of the following joint inventor(s) is missing from the oath or declaration:
Applicant(s) should provide, if possible an oath or declaration signed by the omitted inventor(s), identifying this application by the above Serial Number and Filing Date. A SURCHARGE MUST ALSO BE SUBMITTED AS INDICATED BELOW.

7. The application was filed in a language other than English. Applicant must file a verified English translation of the application and a fee of \$26.00 under 37 CFR 1.17(k), unless this fee has already been paid NO SURCHARGE UNDER 37 CFR 1.16(e) IS REQUIRED FOR THIS ITEM.

8. A \$20.00 processing fee is required for returned checks. (37 CFR 1.21(m)).

9. Your filing receipt was mailed in error because check was returned.

10. Other:
A Serial Number and Filing Date have been assigned to this application. However, to avoid abandonment under 37 CFR 1.53(d), the missing parts and fees identified above in items 1 and 3-6 must be timely provided ALONG WITH THE PAYMENT OF A SURCHARGE OF \$ ~~104.00~~ for large entities or \$ ~~52.00~~ for small entities who have filed a verified statement claiming such status. The surcharge is set forth in 37 CFR 1.16(e). Applicant is given ONE MONTH FROM THE DATE OF THIS LETTER, OR TWO MONTHS FROM THE FILING DATE of this application, WHICHEVER IS LATER, within which to file all missing parts and pay any fees. Extensions of time may be obtained by filing a petition accompanied by the extension fee under the provisions of 37 CFR 1.136(a).

Direct the response to, and any questions about, this notice to the undersigned, Attention: Application Branch.

Raphael Ritz
A copy of this notice **MUST** be returned with response.

For Manager, Application Branch
(703) 557-3831

FORM PTO-1533 (REV. 7-87)

OFFICE COPY

For Office Use Only	
<input type="checkbox"/> 102	<input type="checkbox"/> 202
<input type="checkbox"/> 103	<input type="checkbox"/> 203
<input type="checkbox"/> 104	<input type="checkbox"/> 204
<input type="checkbox"/> 105	<input type="checkbox"/> 205

0070



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: James Raymond Davis and Christopher M. Schmandt
Serial No.: 565,274
Filed: August 9, 1990
For: AUTOMOBILE NAVIGATION SYSTEM

Attention: Application Division

Commissioner of Patents and Trademarks
Washington, D.C. 20231

TRANSMITTAL OF SIGNED DECLARATION AND FEE
PURSUANT TO 37 CFR §§1.53 and 1.16(e)

Sir:

In response to the "Notice to File Missing Parts of Application" mailed August 28, 1990, applicant encloses herewith a Combined Declaration and Power of Attorney, together with a check in the amount of \$527.00 in payment of the required surcharge for filing as a Large Entity. A copy of the Assignment which is being filed under separate cover is also enclosed.

The commissioner may apply any overpayment or charge any additional fees to our Deposit Order Account No. 03-1721. A duplicate copy of this transmittal letter is enclosed to facilitate this process.

Respectfully submitted,



Sam Pasternack
Reg. No. 29,576

Choate, Hall & Stewart
Exchange Place
53 State Street
Boston, MA 02109
(617) 227-5020
September 7, 1990
5035P

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner of Patents and Trademarks, Washington, D.C. 20231, on Sept. 7 1990



0071



UNITED STATES DEPARTMENT OF COMMERCE
Patent and Trademark Office

Address: COMMISSIONER OF PATENTS AND TRADEMARKS
Washington, D.C. 20231

40700 201
120.00 105 PK/HB
H3

SERIAL NUMBER	FILING DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO.
---------------	-------------	-----------------------	------------------

07/26/90 08/09/90 DAVIS

[REDACTED]
SHM PASTERNAK
CHOATE, HALL & STEWART
EXCHANGE PLACE
50 STATE STREET
BOSTON, MA 02110

[REDACTED]

DATE MAILED: 08/29/90

**NOTICE TO FILE MISSING PARTS OF APPLICATION—
FILING DATE GRANTED**

A filing date has been granted to this application. However, the following parts are missing.

If all missing parts are filed within the period set below, the total amount owed by applicant as a large entity, small entity (verified statement filed), is \$ _____.

1. The statutory basic filing fee is: missing, insufficient. Applicant as a large entity, small entity, must submit \$ _____ to complete the basic filing fee and MUST ALSO SUBMIT THE SURCHARGE AS INDICATED BELOW.

2. Additional claim fees of \$ _____ as a large entity, small entity, including any required multiple dependent claim fee, are required. Applicant must submit the additional claim fees or cancel the additional claims for which fees are due. NO SURCHARGE IS REQUIRED FOR THIS ITEM.

3. The oath or declaration:
 is missing.
 does not cover items omitted at the time of execution.

An oath or declaration in compliance with 37 CFR 1.63, identifying the application by the above Serial Number and Filing Date is required. A SURCHARGE MUST ALSO BE SUBMITTED AS INDICATED BELOW.

4. The oath or declaration does not identify the application to which it applies. An oath or declaration in compliance with 37 CFR 1.63 identifying the application by the above Serial Number and Filing Date is required. A SURCHARGE MUST ALSO BE SUBMITTED AS INDICATED BELOW.

5. The signature to the oath or declaration is: missing, a reproduction; by a person other than the inventor or a person qualified under 37 CFR 1.42, 1.43, or 1.47. A properly signed oath or declaration in compliance with 37 CFR 1.63, identifying the application by the above Serial Number and Filing Date is required. A SURCHARGE MUST ALSO BE SUBMITTED AS INDICATED BELOW.

6. The signature of the following joint inventor(s) is missing from the oath or declaration: _____ Applicant(s) should provide, if possible an oath or declaration signed by the omitted inventor(s), identifying this application by the above Serial Number and Filing Date. A SURCHARGE MUST ALSO BE SUBMITTED AS INDICATED BELOW.

7. The application was filed in a language other than English. Applicant must file a verified English translation of the application and a fee of \$26.00 under 37 CFR 1.17(k), unless this fee has already been paid NO SURCHARGE UNDER 37 CFR 1.16(e) IS REQUIRED FOR THIS ITEM.

8. A \$20.00 processing fee is required for returned checks. (37 CFR 1.21(m)).

9. Your filing receipt was mailed in error because check was returned.

10. Other:

A Serial Number and Filing Date have been assigned to this application. However, to ~~APPLICATION DIVISION 40~~ abandonment under 37 CFR 1.53(d), the missing parts and fees identified above in items 1 and 3-6 must be timely provided ALONG WITH THE PAYMENT OF A SURCHARGE OF \$40.00 for large entities or \$55.00 for small entities who have filed a verified statement claiming such status. The surcharge is set forth in 37 CFR 1.16(e). Applicant is given ONE MONTH FROM THE DATE OF THIS LETTER, OR TWO MONTHS FROM THE FILING DATE of this application, WHICHEVER IS LATER, within which to file all missing parts and pay any fees. Extensions of time may be obtained by filing a petition accompanied by the extension fee under the provisions of 37 CFR 1.136(a).

Direct the response to, and any questions about, this notice to the undersigned, Attention: Application Branch.

R. Pasternak
A copy of this notice MUST be returned with response.

040 TL 09/14/90 07565274
For: Manager, Application Branch
(703) 557-5204

FORM PTO-1533 (REV. 7-87)

COPY TO BE RETURNED WITH RESPONSE

1 201	407.00 CK
1 105	
For Office Use Only	
<input type="checkbox"/> 102	<input type="checkbox"/> 202
<input type="checkbox"/> 103	<input type="checkbox"/> 203
<input type="checkbox"/> 104	<input type="checkbox"/> 204
<input type="checkbox"/> 105	<input type="checkbox"/> 205

0072

EXHIBIT

2

(Part 2)

CO-SIGNED DECLARATION AND POWER OF ATTORNEY
#3

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.
I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the

invention entitled AUTOMOBILE NAVIGATION SYSTEM

the specification of which

is attached hereto.

xx was filed on August 9, 1990 as Application
Serial No. 07,565,274 and was amended on _____

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, §1.56(a).

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:
Prior Foreign Application(s):

(Number)	(Country)	(Day/Month/Year Filed)	Yes	No

I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith: Sam Pasternack, Esquire 39,576.

Address all telephone calls to Sam Pasternack, Esq. at telephone no. (617) 227-5020.

Address all correspondence to Sam Pasternack, Choate, Hall & Stewart, Exchange Place, 53 State Street, Boston, MA 02109

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patents issued thereon.

Full name of sole or first inventor James Raymond Davis 40100

Inventor's signature *J. Raymond Davis* Date: 121 AUG 10

Residence: 140 Elm Street, North Cambridge, MA 02140

Citizenship: United States

Post Office Address: 140 Elm Street, North Cambridge, MA 02140

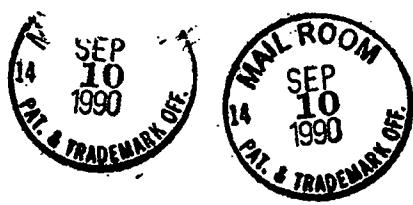
ADDITIONAL INVENTORS

40200
Full name of inventor Christopher M. Schmandt Date: 3/21 Aug 90
Inventor's signature ✓ Ch Eno
Residence 15 Martin Road, Milton, MA
Citizenship United States
Post Office Address 15 Martin Road, Milton, MA

Full name of inventor _____ Date: _____
Inventor's signature _____
Residence _____
Citizenship _____
Post Office Address _____

Full name of inventor _____ Date: _____
Inventor's signature _____
Residence _____
Citizenship _____
Post Office Address _____

Full name of inventor _____ Date: _____
Inventor's signature _____
Residence _____
Citizenship _____
Post Office Address _____



0075

Loc 03B #4 A/N



In the United States Patent and Trademark Office

Applicant : James Raymond Davis and Christopher M. Schmandt Examiner :
 Serial No. : 565,274 Art Unit :
 Filed : August 9, 1990
 For : Automobile Navigation System

Commissioner of Patents and Trademarks
 Washington, D.C. 20231

RECEIVED

SEP 10 1990

Information Disclosure Statement

APPLICATION DIVISION

Sir:

Aspects of the invention have been described in the following sources which were incorporated by reference in the specification and are listed on the attached PTO-1449. Copies are enclosed.

1. "Synthetic speech for real time direction-giving," by C.M. Schmandt and J.R. Davis (*Digest of Technical Papers, International Conference on Consumer Electronics*, Rosemont, Illinois, June 6-9, 1989) is an abstract describing the goals of the research which resulted in the present invention.
2. "Synthetic speech for real time direction-giving," by C.M. Schmandt and J.R. Davis (*IEEE Transactions on Consumer Electronics*, 35(3):649-653, August 1989) is an expansion of the above abstract into a paper. (Kindly note that the publication date for this issue was September 8, 1989; as indicated in the accompanying copy of the certificate of copyright registration.)
3. "The Back Seat Driver: Real time spoken driving instructions," by J.R. Davis and C.M. Schmandt (*Proceedings of the IEEE Vehicle Navigation and Information Systems Conference*, Toronto, Canada, September 1989) describes the strategies employed by the present invention to successfully use speech.
4. "Back Seat Driver: Voice assisted automobile navigation," by J.R. Davis (Ph.D. thesis, Massachusetts Institute of Technology, September 1989) is the most detailed publication describing the present invention to date. The thesis includes a long list of references. Those deemed by the applicants relevant to the present invention as claimed are included on the enclosed PTO-1449 and are discussed below. If the examiner requires further information regarding any of the references cited in the thesis but not included in this Information Disclosure Statement, the applicants will be pleased to provide such.

A short news article on the invention appeared in the July 1990 issue of *Technology Review*. The article, entitled "Terminal Back Seat Driver," is listed on the attached PTO-1449 and a copy of the article is enclosed.

The following references were incorporated and discussed in the specification. They are listed on the attached PTO-1449 and a copy of each is enclosed.

- 1. "CD-ROM Assisted Navigation Systems" by O. Ono, H. Ooe, and M. Sakamoto (*Digest of Technical Papers, IEEE International Conference on Consumer Electronics*, Rosemont, Illinois, June 8-10, 1988) describes the vehicle location system built by NEC Home Electronics, Ltd. which was employed in the working prototypes of the present invention. It is a dead-reckoning position keeping system which uses speed and direction sensors. To compensate for error, it uses map matching on a map database stored on CD-ROM.
- 2. "Attention, intentions, and the structure of discourse" by B.J. Grosz and C.L. Sidner (*Computational Linguistics*, 12(3):175-204, 1986) describes a discourse theory.
- 3. "The intonational structuring of discourse" by J. Hirschberg and J. Pierrehumbert (*Proceedings of the Association for Computational Linguistics*, 136-144, July 1986) describes a theory of intonational meaning.

The following references, listed on the attached PTO-1449, relate to further aspects of the present invention as embodied in the working prototype:

1. *Geographic Base File GBDF/DIME: 1980 Technical Documentation* (U.S. Department of Commerce, Data Users Services Division, 1980) describes the DIME format, the basis for the map database of the working prototypes of the present invention.
2. "A formal basis for the heuristic determination of minimum cost paths" by P.E. Hart et al. (*IEEE Transactions on SSC*, 4:100-107, 1968) is the first of many papers which discuss the A* search algorithm on which the route finder of the working prototypes of the present invention is based.
3. "Cellular telephone data communication system and method" by Harry M. O'Sullivan (U.S. Patent 4,697,281) describes the error correcting modem used in a working prototype of one embodiment of the present invention.
4. "A voice interface to a direction giving program" by James R. Davis (Technical Report 2, MIT Media Laboratory Speech Group, April 1988) describes the method of address entering using a cellular telephone keypad employed in one embodiment of the present invention.

The following discussion, summarized primarily from the Davis thesis cited above, examines prior automobile navigation systems. A related survey can be found in "Automated provision of navigation assistance to drivers" by Matthew McGranagh et al. (*The American Cartographer* 14(2):121-138, 1987). The references cited are included on the enclosed PTO-1449. The undersigned attorney does not currently possess copies of most of these articles, and therefore copies of most are not enclosed. However, at the request of the examiner, copies of any article cited will be obtained and forwarded to the examiner for consideration.

Early application of computers to navigation was intended to reduce traffic congestion by providing route information to drivers. In the Electronic Route Guidance System (ERGS) described in "An electronic route guidance system for highway vehicles," by D.A. Rosen et al. (*IEEE Transactions on Vehicular Technology* VT-19(1):143-152, Feb. 1970), a driver beginning a route finds the intersection closest to the destination, then enters a five letter code word for the intersection. When the vehicle passes over an induction loop sensor in the road, it transmits the destination to a central computer. The computer determines the best route, and relays instructions to the car. This interchange of information occurs at every instrumented intersection. Driving direc-

tions combine simple text from a nine word vocabulary and directional arrows, and are displayed on a "heads-up" display. The ERGS system was designed but never implemented. A similar system was designed and tested in Germany in the late seventies, and is described in "Function, Equipment, and Field Testing of a Route Guidance and Information System for Drivers (ALI)" by P. Braegas (*IEEE Transactions on Vehicular Technology*, VT-29(2):216-225, May 1980).

The pioneering work on computer navigation assistance is presented in "Let your fingers do the driving: Maps, yellow pages, and shortest path algorithms" by R.J. Elliot and M.E. Lesk (Technical Report, unpublished, Bell Laboratories, 1982) and "Route finding in street maps by computers and people" by R. J. Elliot and M.E. Lesk (*Proceedings of the National Conference on Artificial Intelligence*, pp. 258-261, 1982). Elliot and Lesk showed that specialized, as opposed to general purpose, graph search algorithms are desirable for route finding. One reason for this is that the shortest route may not be the best route. Elliot and Lesk added a system of weights, which are extra costs associated with turns. This caused the algorithm to prefer slightly longer routes with fewer turns to short, twisty routes. Elliot and Lesk also were the first to implement a program to generate written natural language driving instructions for the route. In their instructions, a route consists of a beginning, a sequence of turns and crossings (of rivers or railroads), and an ending. For each of these, there is a template comprising a sequence of words and slots, representing fixed and variable components of a sentence for a given type of act. The words are copied directly to the output, and the slots are filled in according to the particulars of an act. A third contribution of Elliot and Lesk was to integrate the digital map with other location oriented databases, including a services database and a personal address book. This allowed the program to find routes to addresses given a person's name, to find the closest store of a specified category, and to mention stores along the route as possible landmarks.

"Direction assistance" by J.R. Davis and T.F. Trobaugh (Technical Report 1, MIT Media Laboratory Speech Group, Dec. 1987) describes a system which provides spoken directions between locations in the Boston area. It uses a Dectalk speech synthesizer. This synthesizer includes a telephone interface, so it can answer a phone call and decode touch tone button presses. To use Direction Assistance, you call it from a touch tone phone. It answers the call, and prompts you to enter your origin and destination locations as street addresses. It finds a route, then describes the route to you. Direction Assistance was directly inspired by Elliot and Lesk, and extends their work in three ways. The most significant difference is that Direction Assistance speaks its directions, where Elliot and Lesk drew maps and provided written text. The disadvantage is that users must remember the instructions or write them down. A second significant difference between Direction Assistance and the work of Elliot and Lesk is that Direction Assistance generates better quality descriptions of the route. The improvement arises because the text generation process first analyzes the route into a sequence of "acts", and then generates descriptions from these acts, instead of working directly from the route. An act represents something that the driver does. There are eleven different acts, each representing a different way of moving. The type of act depends upon topology (how many streets are present at an intersection, and which way traffic can flow), geometry (what angles the streets make) and what kinds of streets are involved. There is also a

function to find an appropriate cue or landmark (e.g. a street crossed or an underpass) just before the location of the act. The Direction Assistance route finder uses a different algorithm (the A* search) than Elliot and Lesk, and has a different set of weights. The weighting scheme ranks roads by a four-valued "goodness" feature and penalizes routes that use less good roads by multiplying the mileage by a constant factor. It also reduces or waives the penalty for turning under a set of circumstance having to do with predicted ease of following; for example, a turn onto a one way street incurs a lesser penalty, since it is unlikely that the driver would turn the wrong way. It reduces the penalty for "T" turns since the driver cannot possibly miss the place to turn. In practice, these weight changes have very little effect.

The Hertz car rental company offers "Computerized Driving Directions" at some of its rental counters. The directions include approximate mileage and estimated travel time, but are highly schematic, even cryptic. Despite appearances, these instructions are not computer generated. The Hertz system is more akin to a database retrieval system than a route finder. A California firm, Navigation Technology, sells a product called "DriverGuide" which is reported to be able to print driving directions between any two points in the San Francisco area ("Mapping out a new idea" by R. Rosenberg, *The Boston Globe*, February 17, 1987, p. 39 and "A Way to Go From A to B" by R. Alexander, *The New York Times*, March 11, 1989, p. 52).

Peeder Ma describes a system which gives textual directions in "An algorithm to generate verbal instructions for vehicle navigation using geographic database" (*The East Lakes Geographer*, 22:44-60, 1987). His work is similar to both Elliot and Lesk's and to Direction Assistance, but was apparently created independently of both. Ma uses A* search with a penalty factor to minimize the number of turns. Unlike Elliot and Lesk, he uses the same penalty for both left and right turns. His street map representation does not include one-way streets or restrictions on turning ("on left turn") so it does not always find usable routes. His route descriptions use a taxonomy about as elaborate as that of Direction Assistance, but the text generated is more stylized.

The systems of Elliot and Lesk, Ma, and Hertz provide static, textual instructions. Direction Assistance gives static verbal instructions. The limitations inherent in static navigation systems were discussed in the specification.

Several groups have built position or navigation systems for use in automobiles. For the most part, these systems have not been well described in the literature, probably from a desire to preserve commercial secrecy. The following discussion summarizes the applicants' knowledge of such systems, as gained through the cited articles.

The most well known automobile navigation system in the United States is the ETAK Navigator, which displays the car's position on a map display on the dashboard. The system is described in "Extending Low Cost Land Navigation Into Systems Information Distribution and Control" by S.K. Honey, M.S. White and W.B. Zavoli (*IEEE Position and Locations Symposium*, pp. 439-444, 1986, IEEE 86CH2365-5) and "Map Matching Augmented Dead Reckoning" by W.B. Zavoli and S.K. Honey (*Proceedings of the 35th IEEE Vehicular Technology Conference*, pp. 359-362, 1986, IEEE CH2308-5). The map rotates as the car turns so that the forward direction is always straight up on the map. The system provides a limited amount of navigation assistance. The driver may enter a destination (as a street address or intersection), and the system will dis-

play the direction and distance to the point. It remains the driver's task to select an appropriate route to the destination.

The Routerechner, described in "On board computer system for navigation, orientation, and route optimization" by P. Haeussermann (Technical Paper Series 840483, Society of Automotive Engineers, 1984), was designed to provide directions in and between German cities. The route finder could receive real time traffic information by digital radio while on route. This system's map included only the Autobahn, and not the cities (this was before CD-ROMs were widely available), yet it also provided a limited navigation service within cities. The user entered the destination as a pair of coordinates, and the system displayed the direction and distance to the destination. As with ETAK, it was the user's responsibility to select an appropriate road.

The Honda Electro Gyro-Cator, described in "Electro gyro-cator: New inertial navigation system for use in automobiles" by K. Tagami et al. (Technical Paper Series 830659, Society of Automotive Engineers, 1983), provided displayed position of the car by plotting a point on a screen. The driver could determine position by placing a transparent map over the screen. This system did not provide route directions.

The Nissan-Hitachi car navigation and information system displays position on a map, finds the shortest route to a destination taking into account real time traffic information, and gives directions by arrows on the face of a display. The system is described in "Navigation systems using gps for vehicles" by T. Itoh, Y. Okada, A. Endoh, and K. Suzuki (Technical Paper Series 861360, Society of Automotive Engineers, 1986). The system also includes a "secretary mode" which displays the driver's appointments. The system uses a CD-ROM for map data, and combines satellite positioning with dead reckoning.

"Eva: An electronic traffic pilot for motorists" by O. Pilsak (Technical Paper Series 860346, Society of Automotive Engineers, 1986) and "Digital Map Data Bases for Autonomous Vehicle Navigation Systems" by E.P. Neukirchner and W. Zechnall (*IEEE Position and Location Symposium*, pp.320-324, 1986, IEEE 86CH2365-5) describe an automobile navigation system which was developed in Germany by Blaupunkt, the University of Karlsruhe, and the federal government. The system accepts destinations as street addresses, finds minimum time routes, and gives directions by a combination of simple (arrow) graphics and voice. The system can recover from a driver error in following the route and find a new route within 50 meters of travel.

The Phillips corporation, in the Netherlands, is developing a prototype car information and navigation system called CARIN, which is described in "Applications of the compact disc in car information and navigation systems" by M.L.G. Thoone and R.M.A.M. Breukers (Technical Paper Series 840156, Society of Automotive Engineers, 1984) and "Digital maps on compact discs" by H.J.G.M. Benning (Technical Paper Series 860125, Society of Automotive Engineers, 1986). With this system, the driver enters a destination using either a keyboard or a touch sensitive screen. The system displays routes on a map and gives spoken driving instructions. The map is stored on board in CD-ROM, and a radio link provides for updates on traffic conditions. Very little has been published about the system. It is not clear whether the CARIN system retains a map as a "vestigial" display, or because its makers do not appreciate the superiority of speech, or because they see a need for positional information other than route finding.

An article in *Automobile Electronic News* entitled "Blaupunkt adds guidance system" (Vol. 1, No. 12, Monday 3 July 1989, pg. 22.) reported that Blaupunkt (a wholly-owned subsidiary of Robert Bosch G.m.b.H.) had recently begun marketing in West Germany a CD-ROM based navigation system called Travelpilot. The system is similar to ETAK. It uses dead-reckoning and displays vehicle destination and position on a map. The map scale and position continuously adjust. There are nine possible views that can be displayed, showing distances that range from 200 meters to 50 kilometers. Currently only one CD is available with digitized maps of West Germany. It contains the street maps of 83 major cities (with populations greater than 100,000) and all towns with populations of 500 or more. It also contains all major highways and roads with the locations of airfields, country roads, and main highways stations. Street and town names are all stored on the CD to enable the driver to enter his intended address via a simplified menu. The article noted that in the future maps will show the locations of hotels, restaurants, and castles. It reported that Blaupunkt planned to add route guidance in the future. The article makes no mention of the use of speech. It was also noted that Hertz car rental company has begun to advertise the availability of Travelpilot on some of its higher-end models and that talks are being held with some European automakers concerning OEM distribution as well.

An article in *Electronic Engineering Times* entitled "Car map system OK'd" (August 7, 1989) reported the award by the British government of a license for a vehicle guidance system. The first license was awarded provisionally to a consortium headed by the General Electric Company for a pilot scheme in central London, expected to be introduced by 1992. If the results of monitoring by the Department of Transport show that a large-scale system would not prejudice road safety or good traffic management, a second license may be negotiated which would lead to a fully commercial system covering an area of about 1000 square miles within the M25, the motorway that circles London. These commercial operations could begin by the end of 1993. It is expected that eventually, continuous guidance could be provided for most of Britain.

The system, called Autoguide, gives drivers recommended routes to their destinations using a simple graphical display fitted to the vehicle's instrument panel. The in-car unit instructs the driver where to turn and which lane to use with a simple system of arrows and other graphical symbols. The article notes that later on, other information, such as the availability of parking spaces, could be added. The article makes no mention of the use of speech. The heart of the system is the central computer which collects journey times from roads in and around the area covered. Preferred routes are continuously updated, and this information is broadcast to subscriber vehicles from a network of strategically located roadside beacons using infrared communications techniques.

The same article notes that Autoguide is one of several approaches to road navigation in Europe. For example, the Common Market is undertaking a research program called Drive (for Dedicated Road Infrastructure for Vehicle Safety in Europe), which involves wide-ranging studies of road transport information. A similar research scheme, called Prometheus, is reportedly being operated by several European car manufacturers under the 19 nation Eureka program. Another Eureka project is Carminat, in which Dutch and French companies are developing a system that will integrate car navigation, communication, and diagnostics. The article further notes that in West Germany,

Robert Bosch G.m.b.H. and Siemens A.G. have demonstrated car navigation systems. One of the West German developments is said to be compatible with Britain's Autoguide, giving the eventual prospect of a road navigation system that could be used throughout the whole of Europe. The article notes that Britain is probably farther ahead than most countries in applying this technology.

An article in *Automotive Electronic News* entitled "U.K. Picks GEC to Head Navigation Project" (Vol. 1, No. 16, August 28, 1989, pg. 31) discusses a follow-on to Autoguide which uses roadside infra-red beacons to transmit information. Drivers reportedly enter destinations on black and white flat-screen displays, and the system provides distance, bearing, and best route. The article states that an audible signal and voice synthesizer instructs the driver when to turn or change lanes, although no details on how this is done are provided. A copy of this article is enclosed.

An article in *Automotive Electronics Journal* (Vol. 2, No. 9, April 23, 1990, pg. 22) reports that Toshiba seeks OEM partners in automaking for its navigation and guidance system. According to a Toshiba spokesman the system will have voice recognition and speech synthesis because "when you're driving, the part of the body that you're not using is the ears and the mouth." The article only discusses the intention of the manufacturer and provides no implementation details.

An article in *The New York Times Magazine* entitled "Softening of the Arteries" (August 26, 1990) discusses an experimental project called Pathfinder which involves a combination of technologies, including Travelpilot, an on-board navigational system manufactured by ETAK. Travelpilot stores map data on a disk and displays it on a screen attached to the dashboard. For the experiment, Travelpilot has been installed in 25 cars. The roadbeds of a 14-mile stretch of the Santa Monica freeway have been fitted with sensors to monitor the speed and density of traffic. This information is sent to a central computer and then relayed to the cars. There, the information is displayed symbolically on the Travelpilot computer screen. Additional data can be supplied in word messages or via digitized voice messages. In this way, the drivers are alerted to the path of least resistance. A copy of this article is enclosed.

The following discussion describes map databases, particularly with regard to features beyond those present in the DIME format.

The TIGER format, which has several improvements over the DIME format, is described in detail in the following sources: "Principal components of the census bureau's TIGER file" by J. Sobel (Research in Contemporary and Applied Geography Discussion Series 3, Department of Geography, SUNY Binghamton, 1986), "The TIGER system: Automating the Geographic Structure of the United States Census" by R.W. Marx (*Government Publications Review*, 13:181-201, 1986), "GIS, TIGER, and Other Useful Acronyms" by R.W. Marx (*National Conference of Geographic Information Systems*, Canadian Institute for Surveying and Mapping, March 1989), "Programs for assuring map quality at the bureau of the census" by R.W. Marx and A.J. Saalfeld (*Fourth Annual Research Conference*, Geography Division, Room 3203-4, Bureau of the Census, U.S. Department of Commerce, Washington DC 20233, March, 1988), "The TIGER Structure" by C. Kinnear (*AUTO CARTO 8 International Symposium on Automation in Cartography*, April 1987), and "Topology in the TIGER file" by G. Boudriault (*AUTO CARTO 8 International Symposium on Automation in Cartography*, April 1987). As

stated in the present specification, the TIGER format has several improvements over the DIME format, but is still a planar graph representing only physical connectivity. It could be used as the basis of a map database for an automobile navigation system if the extensions discussed in the specification are incorporated.

The Taxi! driving simulator, described in "Taxi! Dynamic cartographic software for training cab drivers" by M. Bosworth and R. Low (Technical report, Hunter College Department of Geology and Geography, (212)-772-4000, 1988 paper presented at the Annual Meeting of the Association of American Geographers) includes the concept of turn resistance, a number from one to ten specifying the difficulty of making a given turn, in its map database.

Neukircher and Zechnall, 1986, describes the features of the map used in the Eva system. This map has better position information than DIME. Points are stored in three dimensions and are accurate to 2.5 meters. Road segments are straight lines, chosen so that a new segment begins at either an intersection or when the change in direction exceeds 30 degrees, or when the distance from the center line exceeds 5 meters. Additional attributes of the roads include height and weight restrictions, location of magnetic anomalies, warnings, landmarks, special objects useful in descriptions (e.g. underpasses), layout of complex intersections, and signs. The map has two levels of detail. The coarse level is used for route finding, and the fine level has more detailed information for position finding. Route finding information includes two values for expected speed (one for normal conditions and a second for times of high density), the expected wait time at segment endpoints, and areas where children are likely to be playing.

The University of Calgary AVL-2000 system uses a map that originated as a Canadian government Area Master File. This format, similar to DIME, also required extensive augmentation, as described in "Digital Map Dependent Functions of Automatic Vehicle Location Systems" by C.B. Harris, L.A. Klesh, E.J. Krawkiwshy, H.S. Karimi, N.S.T. Lee (*IEEE Position and Location Symposium*, pp. 79-87, 1988, IEEE CH2675-7). Link (segment) attributes include distance, expected travel time, safety, scenic value, tolls, "impediment value" [sic], one way limitations, banned turns, road type (over- and under-pass, traffic circles, clover leaf), presence of meridians (divided?), and restricted areas. Harris describes as a "special problem" those "source of destination points which correspond to a street addresses [sic] which do not have a unique node identifier". Either their map representation cannot interpolate addresses along segments, or the route finder is restricted to finding routes to nodes only. Harris also mentions auxiliary road information including landmarks, points of interest, emergency services, commercial establishments, weather conditions, traffic flow, and road characteristics and stresses and importance of being able to update the map database over a communications link while driving.

Most systems have expanded the classification of streets. The ETAK map classification is interstate highway, semi-limited access roads and state highways, arterial, collector, light duty roads, alleys or unpaved roads, high speed ramps and low speed ramps. This rich taxonomy is essential to ETAK for choosing which roads to display (lesser roads are suppressed at larger scales to control detail) and in which colors to display them. The Eva system has a two level taxonomy: *rural*, including motorways

and federal highways with separate directional lanes and without intersections, federal highways, roads wider than six meters, roads four to six meters wide, and others, and *urban*, including divided, through, main, side, and restricted.

These maps have some questionable design decisions on the representation of legal restrictions. The ETAK map has no legal topology at all. It is not intended for route finding. The EVA map apparently encodes restrictions on turning by signs, rather than directly in the network. The Calgary map represents legal topology (one ways, banned turns) as a link attribute instead of in the network topology. It may be that the street network represents only physical topology, with the assumption that legal topology will be equivalent to the physical topology unless specially indicated.

The Back Seat Driver appears to be unique in maintaining separate but equal representations for physical and legal topology. These two topologies should be integrated because legal topology is needed for route finding, and physical topology for route description.

Some navigation systems attempt to give warnings about hazardous conditions. In those systems (Eva, Calgary), the hazards (about slope, width, or curves) are encoded explicitly into the map.

Respectfully submitted,



Sam Pasternack
Registration No. 29,576
Attorney for Applicant

Choate, Hall & Stewart
Exchange Place
53 State Street
Boston, MA 02109
(617) 227-5020

1056

卷一 14

***EXAMINER:** Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

0086

***EXAMINER:** Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

0087

				SHOW <u>4</u> of <u>14</u>			
Form PTO-1449 (REV. 8-88) INFORMATION DISCLOSURE CITATION <i>(Use several sheets if necessary)</i>				ATTY. DOCKET NO. MFT5088S SERIAL NO. 565,274			
				APPLICANT James Raymond Davis, et al. FILING DATE August 9, 1990			
				GROUP <u>234</u>			
U.S. PATENT DOCUMENTS							
EXAMINER INITIALS	DOCUMENT NUMBER	DATE	NAME	CLASS	SUBCLASS	FILING DATE IF APPROPRIATE	
FOREIGN PATENT DOCUMENTS							
	DOCUMENT NUMBER	DATE	COUNTRY	CLASS	SUBCLASS	TRANSLATION	
						YES	NO
OTHER DOCUMENTS (Including Author, Title, Date, Pertinent Pages, Etc.)							
		11	"A voice interface to a direction giving program" by James R. Davis, Technical Report 2, MIT Media Laboratory Speech Group, April 1988.				
<i>Gf</i>		12	"Automated provision of navigation assistance to drivers" by Matthew McGranag et al., The American Cartographer 14(2):121-138, 1987.				
		13	"An electronic route guidance system for highway vehicles" by D.A. Rosen et al IEEE Transactions of Vehicular Technology 19(1):143-152, Feb., 1970.				
EXAMINER			DATE CONSIDERED <u>10/31/91</u>				
<small>*EXAMINER: Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.</small>							

0088

Form PTO-1449 (REV. 8-68)				U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE		ATTY. DOCKET NO. MIT5088S	SERIAL NO. 565,274
INFORMATION DISCLOSURE CITATION (Use several sheets if necessary)				APPLICANT James Raymond Davis et al.			
				FILING DATE August 9, 1990	GROUP 234		
U.S. PATENT DOCUMENTS							
EXAMINER INITIAL	DOCUMENT NUMBER	DATE	NAME		CLASS	SUBCLASS	FILING DATE IF APPROPRIATE
FOREIGN PATENT DOCUMENTS							
	DOCUMENT NUMBER	DATE	COUNTRY		CLASS	SUBCLASS	TRANSLATION YES NO
OTHER DOCUMENTS (Including Author, Title, Date, Pertinent Pages, Etc.)							
	14	"Function, Equipment, and Field Testing of a Route Guidance and Information System for Drivers (ALI)" by P. Braegas, IEEE Transactions on Vehicular Technology, Vol 29(2):216-225, May 1980.					
<i>book</i>	15	"Let your fingers do the driving: Maps, yellow pages, and shortest path algorithms" by R.J. Elliot et al., Technical Report, unpublished, Bell Laboratories, 1982.					
	16	"Route finding in street maps by computers and people" by R.J. Elliot et al., Proceedings of the National Conference on Artificial Intelligence, pp. 258-261, 1982.					
EXAMINER <i>Edward Pipala</i>			DATE CONSIDERED <i>10/31/91</i>				
<small>*EXAMINER: Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.</small>							

0089

7 14

Form PTO-1440 (REV. 8-83) INFORMATION DISCLOSURE CITATION (Use several sheets if necessary)				ATTY. DOCKET NO.	SERIAL NO.	
				MIT5088S	565,274	
				APPLICANT	James Raymond Davis et al.	
				FILING DATE	GROUP	
				August 9, 1990	234	
U.S. PATENT DOCUMENTS						
EXAMINER INITIAL	DOCUMENT NUMBER	DATE	NAME		CLASS	SUBCLASS
FOREIGN PATENT DOCUMENTS						
	DOCUMENT NUMBER	DATE	COUNTRY		CLASS	SUBCLASS
OTHER DOCUMENTS (Including Author, Title, Date, Pertinent Pages, Etc.)						
	2D	"An algorithm to generate verbal instructions for vehicle navigation using geographic database", <i>The East Lakes Geographer</i> , 22:44-60, 1987.				
<i>Gf</i>		"Extending Low Cost Land Navigation Into Systems Information Distribution and Control" by S.K. Honey et al, IEEE Position and Locations Symposium, pp. 439-444, 1986, IEEE 86CH2365-5.				
<i>Gf</i>		"Map Matching Augmented Dead Reckoning" by W.B. Zavoli et al, Proceedings of the 35th IEEE Vehicular Technology Conference, pp. 359-444, 1986, IEEE CH2308-5.				
EXAMINER	Edward Pipala			DATE CONSIDERED	10/31/91	
<small>*EXAMINER: Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.</small>						

0091

Form PTO-1449 (REV. 3-83)		U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE		ATTY. DOCKET NO. MIT5088S	SERIAL NO. 565,274	
INFORMATION DISCLOSURE CITATION <i>(Use several sheets if necessary)</i>				APPLICANT James Raymond Davis et al.		
<i>1990</i>				FILING DATE August 9, 1990	GROUP 234	
U.S. PATENT DOCUMENTS						
EXAMINER INITIAL	DOCUMENT NUMBER	DATE	NAME	CLASS	SUBCLASS	
FOREIGN PATENT DOCUMENTS						
	DOCUMENT NUMBER	DATE	COUNTRY	CLASS	SUBCLASS	TRANSLATION YES NO
OTHER DOCUMENTS (Including Author, Title, Date, Pertinent Pages, Etc.)						
<i>gf</i> <i>gf</i> <i>gf</i> <i>gf</i>	<p>"On board computer system for navigation, orientation, and route optimization" By P. Haessermann, Technical Paper Series 840483, Society of Automotive Engineers, 1984</p> <p>"Electro gyro-cator: New inertial navigation system for use in automobiles" by K. Tagami et al., Technical Paper Series 830659, Soc. of Automotive Engineers, 1983</p> <p>"Navigation systems using gps for vehicles" By T. Itoh, et al, Technical Paper Series 861360, Soociety of Automotive Engineers, 1986.</p>					
	EXAMINER	<i>Edward Pipala</i>		DATE CONSIDERED		<i>10/31/91</i>
	<small>*EXAMINER: Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.</small>					

0092

Form PTO-1440 (REV. 2-28-88)		U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE		ATTY. DOCKET NO. MIT5088S	SERIAL NO. 565,274	
				APPLICANT James Raymond Davis et al.		
				FILING DATE August 9, 1990	GROUP 234	
U.S. PATENT DOCUMENTS						
EXAMINER INITIAL	DOCUMENT NUMBER	DATE	NAME	CLASS	SUBCLASS	
FOREIGN PATENT DOCUMENTS						
	DOCUMENT NUMBER	DATE	COUNTRY	CLASS	SUBCLASS	TRANSLATION YES NO
OTHER DOCUMENTS (Including Author, Title, Date, Pertinent Pages, Etc.)						
	<p>"Eva: An electronic traffic pilot for motorists" by O. Pilsak, Technical Papers Series 860346, Society of Automotive Engineers, 1986.</p> <p>"Digital Map Data Bases for Autonomous Vehicle Navigation Systems" by E.P. Neukirchner et al., IEEE Position and Location Symposium, pp. 320-324, 1986, IEEE 86CH2365-5.</p> <p>"Applications of the compact disc in car information and navigation systems" by M.L.G. Thoone et al., Technical Papers Series 840156, Society of Automotive Engineers, 1984.</p>					
	EXAMINER		DATE CONSIDERED			
			10/31/91			
	*EXAMINER: Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.					

0093

***EXAMINER:** Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

0094

Form PTO-1440 (8-88)		U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE		ATTY. DOCKET NO. MITS088S	SERIAL NO. 565,274
SEP INFORMATION DISCLOSURE CITATION <i>(Use several sheets if necessary)</i>		APPLICANT James Raymond Davis et al.			
		FILING DATE August 9, 1990	GROUP 234		
U.S. PATENT DOCUMENTS					
EXAMINER INITIAL	DOCUMENT NUMBER	DATE	NAME	CLASS	SUBCLASS
FOREIGN PATENT DOCUMENTS					
	DOCUMENT NUMBER	DATE	COUNTRY	CLASS	SUBCLASS
OTHER DOCUMENTS (Including Author, Title, Date, Pertinent Pages, Etc.)					
	35	"The TIGER system: Automating the Geographic Structure of the United States Census" by R.W. Marx, Government Publications Review, 13:181-201, 1986.			
	36	"GIS, TIGER, and Other Useful Acronyms" by R.W. Marx, National Conference of Geographic Information Systems, Canadian Institute for Surveying and Mapping, March 1989.			
	37	"Programs for assuring map quality at the bureau of the census" by R.W. Marx et al., Fourth Annual Research Conference, Geography Division, Room 3203-4, Bureau of the Census, U.S. Department of Commerce, Wash., D.C. 20233, March 1989.			
EXAMINER	Edward Pipala		DATE CONSIDERED	10/31/91	

***EXAMINER:** Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

0096

Form PTO-1440 (Rev. 8-8-83)		U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE		ATTY. DOCKET NO. MIT5088S	SERIAL NO. 565,274
INFORMATION DISCLOSURE CITATION <i>(Use several sheets if necessary)</i>				APPLICANT James Raymond Davis	
12 1990				FILING DATE August 9, 1990	GROUP 234
U.S. PATENT DOCUMENTS					
EXAMINER INITIAL	DOCUMENT NUMBER	DATE	NAME	CLASS	SUBCLASS
					<i>FILING DATE IF APPROPRIATE</i>
FOREIGN PATENT DOCUMENTS					
	DOCUMENT NUMBER	DATE	COUNTRY	CLASS	SUBCLASS
					<i>TRANSLATION YES NO</i>
OTHER DOCUMENTS (Including Author, Title, Date, Pertinent Pages, Etc.)					
<i>gf</i>	"The TIGER Structure" by C. Kinnear, AUTO CARTO 8 International Symposium on Automation in Cartography, April, 1987.				
<i>gf</i>	"Taxi! Dynamic cartographic software for training cab drivers" by M. Bosworth et al, Technical report, Hunter College Department of Geology and Geography, (212)-772-4000, 1988 paper presented at the Annual Meeting of the Association of American Geographers.				
<i>gf</i>	"Digital Map Dependent Functions of Automatic Vehicle Location Systems" by C.B. Harris et al., IEEE Position and Location Symposium, pp. 79-87, 1988, IEEE CH2675-7.				
EXAMINER	<i>Edward Pipalla</i>			DATE CONSIDERED	<i>10/31/91</i>
<small>*EXAMINER: Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.</small>					

0097

Form PTO-1440 (Rev. 6-28-83)		U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE		ATTY. DOCKET NO. MIT5088S	SERIAL NO. 565,274
 INFORMATION DISCLOSURE CITATION (Use several sheets if necessary)				APPLICANT James Raymond Davis et al.	
				FILING DATE August 9, 1990	GROUP 234
U.S. PATENT DOCUMENTS					
EXAMINER INITIAL	DOCUMENT NUMBER	DATE	NAME	CLASS	SUBCLASS
FILING DATE IF APPROPRIATE					
FOREIGN PATENT DOCUMENTS					
	DOCUMENT NUMBER	DATE	COUNTRY	CLASS	SUBCLASS
TRANSLATION YES NO					
OTHER DOCUMENTS (Including Author, Title, Date, Pertinent Pages, Etc.)					
<i>gf</i>	"Softening of the Arteries", By Bruce Weber, The New York Times Magazine, p. 78, August 26, 1990.				
	"Terminal Back Seat Driver", Technology Review, July, 1990, p. 10.				
<i>gf</i>					
EXAMINER	<i>Edward Pipalk</i>		DATE CONSIDERED	<i>10/31/91</i>	
<small>*EXAMINER: Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.</small>					

0098

⑨
**Back Seat Driver: voice assisted automobile
navigation**

by

James Raymond Davis

B.S.A.D., Massachusetts Institute of Technology (1977)

Submitted to the Media Arts and Sciences Section
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1989

©Massachusetts Institute of Technology 1989

All Rights Reserved

Signature of Author.

James R. Davis
Media Arts and Sciences Section
August 4, 1989

Certified by

3 - 2
Nicholas P. Negroponte
Professor of Media Technology
Thesis Supervisor

Accepted by

Stephen A. Benton
Stephen A. Benton
Chairman, Departmental Committee on Graduate Students

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

FEB 27 1990

SEARCHED *SEARCHED*

0111



**Back Seat Driver: voice assisted automobile
navigation**

by

James Raymond Davis

B.S.A.D., Massachusetts Institute of Technology (1977)

Submitted to the Media Arts and Sciences Section
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1989

©Massachusetts Institute of Technology 1989.

All Rights Reserved

Signature of Author **Media Arts and Sciences Section** **August 4, 1989**

Certified by **Nicholas P. Negroponte**
Professor of Media Technology
Thesis Supervisor

Accepted by **Stephen A. Benton**
Chairman, Departmental Committee on Graduate Students

Back Seat Driver: voice assisted automobile navigation

by

James Raymond Davis

Submitted to the Media Arts and Sciences Section
on August 4, 1989, in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

Abstract

The Back Seat Driver is a computer navigation assistant for drivers in a city. It differs from earlier navigation programs by using speech, rather than graphics, to give instructions. The advantages of speech are that the driver's eyes are left free for driving and that the spoken directions contain information not easily portrayed in pictures. The program talks about the features of the road in the same way the driver sees them, giving the impression that the program is actually in the car.

Driving instructions are modeled after those given by people. The two issues for spoken directions are *what to say* (content) and *when to say it* (timing). The content of the instructions tells the driver what to do and where to do it. The program has a large taxonomy of intersection types, and chooses verbs to indicate the kind of intersection and the way of moving through it. The instructions refer to landmarks and timing to tell the driver when to act.

Timing is critical because speech is transient. Drivers hear instructions just in time to take the required action, and thus need not remember the instruction or exert effort looking for the place to act. The program also gives instructions in advance, if time allows, and the driver may request additional instructions at any time. If the driver makes a mistake the program describes the mistake, without casting blame, then finds a new route from the current location.

Street map databases for navigation programs must distinguish between *physical* connectivity (how pieces of pavement connect) and *legal* connectivity (whether one can legally drive onto a physically connected piece of pavement). Legal connectivity is essential for route finding, and physical connectivity for describing the route. The database must also contain all landmark information, since the program has no "eyes".

EXHIBIT

2

(Part 3)

The Back Seat Driver is an actual working prototype. It has successfully guided drivers unfamiliar with Cambridge to their destinations. Although much work remains, it is easy to foresee a practical implementation in the future.

Thesis Supervisor: Nicholas P. Negroponte
Title: Professor of Media Technology

Acknowledgments

I owe gratitude to many people for help with this work. In particular I would like to thank: Phil Rittmueller and NEC Home Electronics (USA) Inc., who sponsored it; Symbolics Inc., Digital Equipment Corporation, the Nippon Telephone and Telegraph company, and DARPA, which sponsored my earlier graduate studies; Brewster Kahle, who provided the original inspiration for my work on computerized direction giving in the summer of 1985; Thinking Machine Corporation, where the work was conducted; and Tom Trobaugh, who began that study with me. I thank those who provided me with helpful information: James A. O'Connell, Jr. and Peter Naghavi of the Somerville Department of Traffic and Parking, Jane Kent of the Cambridge Department of Traffic, and George Hawat of the Boston Transportation Department, for providing traffic light data bases; Doris Walter of the Massachusetts District Commission for maps of rotaries; Doug Milliken for explanations of car dynamics; Don Cooke of Geographic Data Technology and Joel Sobel of the Census Bureau for information on the DIME and TIGER map formats; David Pietraszewski of the United States Coast Guard and Michael Fisher of Trimble Navigation for information about GPS. The photo on page 60 was taken by Kyle G. Peltonen, and is reprinted by permission of The Tech, the Mr. Boffo cartoon is reprinted courtesy of the Tribune Media Services. I thank also Weng-yew Ko for building hardware interfaces; Elaine McNair at NEC for removing many obstacles; Michael Sullivan of Spectrum Communication for solving communications problems; Greg Parro of Cellular Phone Services for excellent service with the tele-

phones; Hidehiro Matsumoto of NEC for translations to and from the Japanese; Janette Noss and Elaine McCarthy, who keep the Architecture Machine Group running; Anna Korteweg, Walter Bender, Gary Drescher, and Carol Strohecker for careful editing, and Shawn T. Williams and Greg Grove, who helped by participating in UROP work. I wish to thank all those who taught me of computational linguistics; especially Barbara Grosz, Julia Hirschberg, and Janet Pierrehumbert. They of course bear no responsibility for my failure to fully grasp the material they attempted to teach me. My committee members were Steven Benton, Mike Lesk, Nicholas Negroponte, and Chris Schmandt. None of these people, of course, are responsible in any way for the ways I have misunderstood their ideas or abused their work, or for my failure to heed their advice. Finally I must specially mention Nicholas Negroponte, who created the Media Lab where this work was possible, and Chris Schmandt, the director of the Speech Research Group, who has been directly involved in everything I have done here at the Media Lab. A full list of his contributions to my studies would require many pages.

Those I have forgotten or otherwise neglected, I can only ask to deepen my debt by forgiving my shoddy memory and inexpressive words.

This thesis is dedicated to my son, Adam, may he find his way back soon.

Contents

1. Overview	16
I. Natural Directions	21
2. Human Direction-Giving	22
2.1 Instructions	24
2.1.1 Verbs	25
2.1.2 Direction	28
2.2 Distance and Time	29
2.2.1 Landmarks	31
2.2.2 Advance Notice	32
2.3 Advice	33
2.4 Style and Packaging	33
2.5 Silence	34
2.6 Example Instructions	36

II. Implementation of the Back Seat Driver	41
3. Cartography	42
3.1 DIME format	42
3.2 The limits of DIME	44
3.2.1 The model of connection is deficient	44
3.2.2 Position resolution is inadequate	45
3.2.3 The forthcoming TIGER format is better, but not enough	46
3.3 A better map	48
3.3.1 Legal connectivity	48
3.4 Additional segment attributes	48
3.4.1 Quality combines size and speed of a segment	49
3.4.2 Expanded street classifications	49
3.4.3 Divided Roads	52
3.4.4 Landmarks help directions	52
3.4.5 Lane information	54
3.4.6 DIME files need correction	54
3.5 Other maps	55
3.6 Improvements are still needed	58
3.6.1 Summary	61
4. Descriptions	62
4.1 Classifying movement	63

4.2 Definitions of actions	65
4.3 What to do	69
4.4 When to do it	71
4.4.1 Timing	71
4.4.2 Cues	72
4.4.3 Landmarks	74
4.4.4 Street names as landmarks	75
4.5 Advice	77
4.6 Discourse	78
4.7 Mistakes	80
4.8 Reassuring	82
5 User and System Goals	83
5.1 Resource allocation	84
5.2 Architecture	87
5.3 Kinds of Goals	88
5.3.1 Finding Services	89
5.3.2 Other user goals	89
5.3.3 System goals	92
6 Comparing Routes	92
6.1 Comparing routes requires a metric	95
6.2 Estimating the time required to find a route	95

III Conclusions	97
7 Related Work	98
7.1 Early Work	98
7.2 Elliot and Lesk	99
7.3 Direction Assistance	101
7.3.1 Entering addresses	102
7.3.2 Generating text	103
7.3.3 Route Finding differences	104
7.3.4 Text-Based Directions	105
7.4.1 Counter Top Directions	105
7.4.2 Maps	105
7.5 Automotive Navigation Systems	106
7.6 Classifying navigation systems	107
8 Future Developments	110
8.1 Integration with the car	110
8.2 Knowing the driver	111
8.2.1 Improving the directions	112
8.2.2 Learning about the driver	112
8.3 Confidence	114
8.4 Speech interface issues	116
8.4.1 Speech output	116

8.4.2. Speech input	117
8.5 Understanding the route	117
8.6 Other directions	120
8.7 Changes to the map	120
8.8 Integration with the city	121
8.9 Policy	122
8.9.1 Liability	122
8.9.2 Privacy	123
9 Summary	125
A Location Systems	127
A.1 Accuracy requirements of the Back Seat Driver	127
A.2 Classifying location systems	128
A.3 Position finding	129
A.3.1 LORAN-C	129
A.3.2 Polled pulse time ranging	131
A.3.3 GPS	132
A.3.4 TRANSIT	133
A.3.5 Other satellite systems	134
A.3.6 Beacons	135
A.3.7 The cellular phone system might serve as a location system	136
A.3.8 Vision	137

A.4 Position keeping	137
A.4.1 Dead reckoning systems need correction	138
A.5 Hybrid systems	140
B Route Finding	142
B.1 Breadth-first search considers all possible partial solutions in parallel	143
B.2 Best-first search saves effort	145
B.3 A* search avoids falsely promising paths	146
B.3.1 A suboptimal, but faster algorithm, is desirable	147
C Communication with the car	149
C.0.2 Cellular phones are hostile for data transmission	150
C.0.3 Retransmission introduces latency	152

List of Figures

2-1	Fulkerson turns into Binney	27
2-2	Map for sample route	40
3-1	Position roundoff makes headings uncertain	45
3-2	Short segments have less angular resolution	46
3-3	Coordinate rounding makes streets seem to wobble	47
3-4	The "super" streets of Boston	50
3-5	A rotary	51
3-6	Access ramps at an interchange	52
3-7	When trucks ignore height restrictions	60
4-1	T, fork, and exit all have same topology	64
4-2	Map near MIT	70
5-1	Resource allocation	85
6-1	Four routes from 12 Albany Street to + Glenville Terrace	93
7-1	A U turn	104

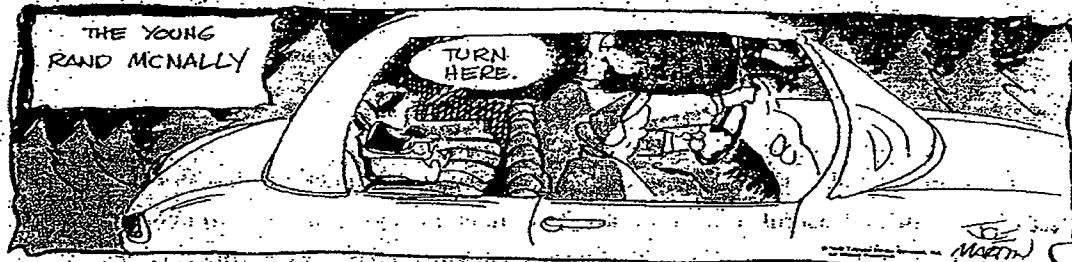
7-2 Example of driving instructions provided by Hertz	105
A-1 Fraction of segments shorter than X in meters	128
A-2 Geometric Dilution of Precision	130
A-3 Dead reckoning errors increase with distance	139
A-4 Map matching corrects dead reckoning errors	140
A-5 Dead Reckoning corrects for loss of GPS signal	141
B-1 Breadth-first search example	145
C-1 Communications block diagram	150
C-2 Histogram of durations of odometer sequence	151
C-3 Probability of cellular call termination increases with time	152
C-4 Histogram of inter-arrival times of packets	153

List of Tables

2.1	Verbs in descending order of frequency	25
2.2	Notation system for transcriptions	36
B.1	Breadth-first search considers many potential routes	146
B.2	A* search touches fewer segments, and is therefore faster	147
B.3	Comparative search times and route lengths for 30 routes with different values of distance weight factor	148

frontispiece

MISTER BOFFO by Joe Martin



Chapter 1

Overview

This thesis is about the design and construction of a machine that does something difficult and useful in a new way. The machine's purpose is help people find their way by car from one place to another within a city, a task which is clearly useful and also worth improvement. A study done for the United States Federal Highway Administration estimated that 45 billion dollars are wasted each year in the U.S. because of ineffective routing, from causes including being lost, stuck in traffic, or choosing bad routes[1]. The machine discussed here is called the Back Seat Driver¹. The Back Seat Driver is a computer program which uses synthetic speech to give instructions to the driver of a car as needed while driving. The Back Seat Driver differs from previous navigation assistance programs by using speech to give directions², instead of drawing a map or displaying symbols.

My two concerns are to determine the best form and content for spoken instructions, and to determine what information a program requires in its map to find routes and provide excellent instructions. My approach towards both questions is

¹The usual sense of this term is an unwanted critic of one's driving skills. This is not what I intend.

²Although there are reports of earlier navigation systems using speech, none are described in the literature.

empirical. The initial design of the program was based on a short study of natural direction giving. This initial design was the "initial conditions" for an iterative process of design, testing with drivers, and revision. When the design was wrong, drivers complained, or just got lost, and I changed the design until they stopped complaining. This thesis describes the design that emerged from this process.

The thesis is divided into three parts. The first part, **Natural Directions**, describes the strategies and styles than people use when giving directions. The second part, **Implementation**, describes the program. Chapter 3 describes the map database. The next chapter tells how the Back Seat Driver decides what to say and when to say it. This is the key to the entire thesis. Chapter 5 tells about the other things the program does (such as read mail), and how it decides which task to work on at a given time. Chapter 6 tells how the Back Seat Driver evaluates routes when searching for the best one. The final part, **Conclusion**, describes related work, then tells how the Back Seat Driver might be improved and suggests directions for future research.

Three appendices provide background material and certain details of the implementation. The first two set out the technologies and concepts used in this program and in vehicle navigation in general. Depending on what you know already, you may have to read one or both of these section before reading earlier sections. The third appendix describes data communications using cellular telephones.

Why speech?

The first issue in design of a vehicle navigation system is the choice of modality for giving instructions. Two channels seem possible: vision and hearing². Given the technical ability to display a map on the dashboard, or to display directional

²We could also consider the haptic channel: This is how people tell horses which way to go. It is unlikely to be popular with users, however.

arrows on the windshield, or the ability to synthesize speech, which should we choose?

Previous in-car navigation systems have used the visual channel to give navigation information. The Back Seat Driver uses speech because there are reasons not to use vision, and some advantages to speech.

Much work has been devoted to developing systems that display the vehicle's position on a map in the car (or at a dispatching facility). When combined with the ability to display a route (or at least the relative or absolute position of the destination), such displays can be used for navigation. One argument against such displays is that they require that the driver look at them while driving, and this makes driving less safe. Although drivers can spare some visual attention while driving along straight roads[88], it is not as clear that they can afford to look at the map while turning, yet it is just while turning that drivers are most in need of navigation information. Visual displays are most easily used when they are least needed.

A second argument against the use of maps is that many people have difficulty finding and following routes on paper maps[81]. That is, they are not "map literate". Information presented on maps will simply be unintelligible to such people. Paper maps do not really qualify as navigation assistance, since they do not show the driver's current position, though they may show the route.

An experiment comparing navigation aids was conducted by Streeter [82], who compared the performance of drivers under three different conditions. One set of drivers received customized paper maps with the route to be driven highlighted in red. A second set of drivers heard spoken instructions from a tape player that permitted them to play the next or previous instruction. Drivers using the maps took longer to arrive than those who had verbal instructions. They also made more errors and drove further. This can be explained by the extra effort required to consult the paper map. The third set of drivers got both navigational aids.

Surprisingly, drivers with *both* sources of information did *worse* than those using only voice, though better than those with a map only.

Neither navigation aid in the Streeter experiment included information about the current position. Both required the driver to determine when to carry out the instruction and to decide whether the instruction was correctly executed. That is, the driver could play an instruction as often as she liked, but had to decide for herself when to advance to the next instruction. The experiment did not compare voice to either an electronic map (indicating current position as well as route) or directional symbols (indicating which way to go). For this reason, the experiment does not directly answer the question about choice of modality. It is, at best, suggestive.

One advantage of speech is that the driver's eyes are left free for driving. In addition, speech uses words, and can therefore refer to past and future actions and objects not yet seen. This is hard to do with symbolic displays or maps.

On the other hand, speech has some problems as well. Speech is transient. The driver must remember what the program said, and the program must be prepared to repeat itself on demand. A second consequence of the ephemeral nature of speech is that the driver has no evidence of the program's operation except when it speaks. In a period of long silence, the driver may fear that the program has failed⁴. The program must do extra work to keep the driver confident in its continued operation. This turns out to be beneficial. Some of the remarks the program makes have little to do with route following per se, but rather are descriptions of the immediate vicinity of the road. When these are uttered at the right time, the driver gets a very strong sense that the program is seeing the world in the same way she does, and this is very reassuring.

The arguments here against maps and for speech convinced me that the Back

⁴Programs - especially prototypes - are not as reliable as, say, telephones. No one worries that a silent phone may be broken. A program silent for too long, though, is cause for concern.

Seat Driver should speak its directions. The next section tells how people speak their directions.

Natural Directions

Part I

Natural Directions

Chapter 2

Human Direction Giving

My investigations began by studying how people give directions when they are passengers in the car. The intent was to discover common patterns in the directions that could be duplicated by the Back Seat Driver. The number of subjects (six) was far too small to justify any general conclusions about direction giving, but that was not the intention. I studied human direction giving to get a starting point for the design of the Back Seat Driver, not an ending point. I built the Back Seat Driver through an iterative process of design, test, and modification which converged to the system described in this thesis. All I needed was a good first approximation, and a sense of the kinds of variation in direction giving style, and for this I think six subjects was sufficient.

Procedure

My six subjects told me how to drive to destinations of their choosing while riding in the car. All were experienced drivers. The subjects actually wanted to

go to their destination, so they had an incentive to give good directions. (They did not all provide the best routes, but route finding was not a subject of this investigation.) Most, but not all, of their utterances were spontaneous. In some cases I asked questions - "What now?" or "How long will I be on this road?". Our conversations were recorded on a cassette recorder in the car. All talk relevant to the route was transcribed, but most personal conversation was not transcribed. The transcriptions included approximate position of the car, as determined from memory. A sample transcription appears at the end of this chapter.

There are limits to the usefulness of this method. By its very nature, there is no way to control for the route or the destination. No two trips led to the same destination, though several had common beginnings. I assume that there is nothing special about the routes, so I can generalize. The usefulness of this data depends upon the assumption that the subjects were competent to give instructions and that the instructions were in the same form that subjects would have wanted to hear, had they been the driver. There is certainly some doubt about this point. We could do better than to emulate subjects who hesitated, gave misleading directions, or simply pointed¹. Another limitation is that the transcriptions show position only coarsely, and velocity not at all, so they can not be used to answer questions of timing.

Overview

My understanding of driving instructions comes from treating driving as decision making. I think of the driver as constantly aiming the car at a moving target, a patch of pavement, some few yards ahead. The driver is making a new decision several times a second. Some of these decisions involve choosing the next street to go along. When following instructions, at each point of decision one of two con-

¹In fact, some navigation systems do no more than point in the right direction.

ditions must apply; either the instructions must say explicitly what to do, or the driver and instruction-giver must mutually believe that the thing to do is obvious and does not need to be said. This is by far the more common condition. It is only at an intersection that there is any meaningful choice, and even there there is a presumption that drivers will continue forward. Giving instructions should be the exception, not the rule. The traditional "back seat driver" - the annoying critic, not the program I built - bothers the driver by giving frequent warnings about conditions the driver is perfectly aware of, acting as if the driver had no common sense.

Navigators must give the driver the information she needs at the time she needs it. At the most basic level, this information consists of instructions. Instructions concern what to do and when (or where) to do it. People also provide advice which concerns unseen conditions ahead that the instruction giver can anticipate based on experience. The driver can ignore advice, but it may then be harder to follow the instructions. There is also orientation information, describing the surroundings, and placing the route in a larger image of the city. Instructions are essential; advice helpful; and orientation optional. Although I think orientation is important for future investigation, I do not address it in this study of natural direction giving.

2.1 Instructions

In the instructions I recorded, people used many different verbs to describe motion through intersections. I believe that people choose particular verbs to help describe the shape of the intersection and the kind of movement through it. One object of this study is to identify the reasons people prefer one verb to another, and the syntactic constructions in which verbs are employed. In addition to the kind of motion, instructions must also specify the direction of motion. This they

do either with the words "left" and "right", or by naming landmarks in the desired direction, or both.

Now let us examine some of the verbs people use.

2.1.1 Verbs

People use a variety of verbs to describe motion through the streets. Table 2.1 lists the ten most common verbs used in directions, in descending order of frequency in my data.

verb	count	verb	count
take	35	turn	14
bear	22	follow	4
go	22	make	3
keep	13	get	3
stay	8	continue	2

Table 2.1: Verbs in descending order of frequency.

On what grounds do people choose one verb over another? It would be naive to expect to find a single meaning for each verb. The verb "go" is used in several different contexts: "go straight" and "keep going" and "go right". There is almost no situation where the verb "go" is not used. Fortunately, some other verbs have more restrictive contexts. In this study, I concentrate on the more specific verbs to the exclusion of the more "generic" ones, on the grounds that automated directions should use the most specific verb that is still correct. The goal is not to reproduce natural speech, with all its complexity. I am not trying to simulate human behavior; I am trying to impart information concisely.

The verb "take" designates a turn: To turn is to change heading by more than (about) 45 degrees, at an intersection where there are always at least two possible

ways to go (though not necessarily both legal). There are two sizes of turns; "hard" turns (more than 90 degrees) and ordinary turns. After a turn, the car is on a "different" street than it was before.

Although the data does not show when two streets are the "same" and when they are "different", it does let us rule out some possible definitions. For instance, it is clear that change of name alone is not sufficient to make a street "different".

There are plenty of streets that "change names" at an intersection. For example, if one drives up Ames Street from the Media Laboratory, and crosses Main Street in Kendall Square, the street name changes to Sixth Street. Certainly no one would call this a "turn", and nobody would say they were on a different street, either. Other examples are Hampshire Street in Cambridge, which is called Beacon Street in Somerville; and the O'Brien Highway in Cambridge, which is the McGrath Highway in Somerville (and which, by the way, I have never heard called anything other than McGrath). Aside from these counter-examples, name change can not possibly be the criteria for "different", because many drivers do not even know the names of the streets².

Change of heading is also not sufficient to make a turn. Near MIT there is an intersection where Fulkerson Street "turns" into Binney Street. An illustration of this place appears in figure 2-1. A driver proceeding up Fulkerson from Main has no choice about which way to go because the street is divided to prevent her from either continuing up Fulkerson or turning left onto Binney. Nobody calls this a turn, even though both angle and name change. Instead, the road seems to be just curving around to the right. This intersection is, in a way, a pun on the word "turn", which, in the phrase "Fulkerson turns into Binney", means "becomes". On the other hand, near Harvard Square there is a place where Appleton Street makes a right hand turn, and the subject whose route passed through this intersection described it as a turn, though the name was the same.

²But it will turn out that that the program does have to use names as part of its concept of "different", for reasons explained later.

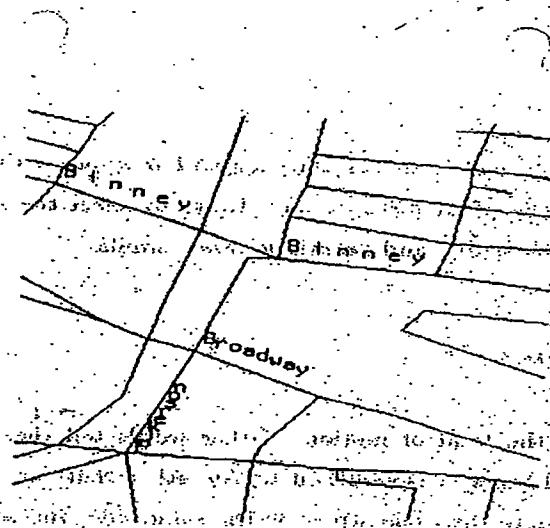


Figure 2-1: Fulkerson turns into Binney

A necessary feature for a "turn" is that there be choice about which way to go, even if the other choice is illegal.

The usual syntactic form for "take" is "take a DIRECTION", where DIRECTION is either "left" or "right". Other verbs with the same meaning as "take", but used less often, are "make", "hang", and "turn". The first two have exactly the same syntax as "take".

To "bear" is to make a lesser change of heading than a turn. After bearing, one may or may not be on a different road. The most common construction is "bear DIRECTION", with variant forms "bear to the DIRECTION" and "bear off to the DIRECTION". In order to "bear" the change in heading must be small and there must also be at least one other road which also requires only a small change in heading. Nobody used the word "bear" to describe the move from Albany Street to Main Street (a change of about 30 degrees). So using the word "bear" not only tells the driver how much turning will be required, it tells something about how the intersection will look. There is an exception. In one case, a subject used the word "bear" to describe motion when there was no choice at all, saying "bear left as the road goes around".

The main use of "keep" is to continue forward motion on the same street³. When subjects used "keep" at forks, it was always to select the straighter of the two alternatives. Verbs "stay" and "continue" were similar.

2.1.2. Direction

Verbs tell about the kind of motion. Other words tell the direction of the motion. In almost all cases, it is sufficient to say either "left" or "right". In some cases, there will be more than one street on the same side, and some other tactic is required.

One is to name a landmark that lies in that direction, e.g. "you bear to the left here and go under the bridge" or "straight to where those lights are." This tactic can also be used when there is no ambiguity, and has the advantage that it also works for people who confuse left and right. In addition, combining a landmark with a direction name ("Go left towards the blinking lights.") adds redundancy to the instructions, thus making them easier to follow. Redundancy is a feature worth incorporating in the Back Seat Driver.

Other approaches to ambiguity are less successful. One subject tried to specify the street using "clock face" terms, where 12 o'clock is straight ahead, 3 o'clock is right, and 9 o'clock is left: "bear to the left about ten o'clock" but when this did not work, he simply pointed. Another ordered the streets by amount of turning: "Now here you want to bear left. Not all the way left but straight to where those lights are". Note that this specification of directions also uses a landmark.

³This surprised me. In my own speech it is synonymous with "bear".

2.2 Distance and Time

In addition to knowing what to do, the driver must know where (or when) to do it. Notably absent from the directions are units of distance. Subjects *could* have given instructions by saying "Drive 310 yards; then ...", but did not. When subjects did give distances, they used qualitative terms like "in a while", "up ahead", "soon", or "not too long after ...". This is not simply because subjects are ignorant of the distances. When I asked them how long we would remain on the current road, most answered with units of either miles or minutes, and the answers were reasonably accurate. Subjects did have some idea of distance, but instead chose other means to tell the driver when to act.

The most common strategy is to give the instruction immediately before the action, just early enough that the driver has time to slow down for the turn.⁴ Instructions given immediately often include a reference to "here" or "right here" as the place to act.

I think that subjects use this form in part because they are not able to formulate the instruction until they are actually at the place of choice. In Benjamin Kuipers' model of navigation ([46], [47]), and discussed in 3.5) the earliest form of route knowledge is a "felt path" where a route is a sequence of pairs of scene description and action. The navigator can not give the next instruction until actually at the place described. Suggestive evidence for this account is that some drivers broke up instructions into two parts, first telling how to get to the next choice point, and then giving the next instruction only when at the point, as shown in this excerpt from a transcription⁵:

⁴This is an assumption, not a fact. I did not measure the relationship between vehicle speed, distance to the intersection, and the time when the passengers spoke. Nor do I have evidence that passengers could reliably estimate these quantities. The subjects had all driven for at least five years. It would be instructive to hear the instructions given by those who have never driven.

⁵The transcription conventions are explained at length below. Briefly, punctuation represents features of spoken delivery, not grammatical form, so sentences do not generally end with periods.

Just keep going until () I think it's until the next major intersection
at least it's the next light anyway

Take a left at the light

Further evidence is that subjects sometimes changed their instructions when approaching the intersection:

Now we're gonna pass Harvard Ave and it's gonna be the next (no)
maybe we're gonna turn onto Harvard Ave come to think of it gonna
turn right () on Harvard

To make sense of this, assume that people are "giving instructions" according to a deficient mental map. When they see the actual intersection, they are able to recall the correct instructions.

Subjects who do have an accurate mental map can give the instruction at the earliest moment when it is unambiguous, that is, when the turn is the next turn, no matter how far away it is. They can also count ahead, and express the distance in units of blocks, turns, or lights ("take the second right"). This gives the driver plenty of notice without requiring much extra work. But such a count can be ambiguous if some of the objects counted are uncertain instances of the category. A small alley or a blinking light might or might not be included in the count. Either for this reason, or lack of knowledge, subjects did not make much use of counting.

One place where it is useful to give distance by a measure is in going around rotaries. In the one case where a route included a rotary, the subject said "You're gonna go around three quarters of the way and head across the bridge". In a rotary, the only appropriate units are those of angular distance, since there may not be signs, and the exits come up quickly.

Second, the empty parenthesis "()" represents a brief pause.

2.2.1 Landmarks

Subjects can tell the driver when to act by naming a landmark at the place of action. Perhaps the best example of use of landmarks is in an instruction for the turn from Massachusetts Avenue onto Back Street, just after crossing the Harvard Bridge into Back Bay, Boston:

At the very end of the bridge here there's um a v- hard right which is h- very hard to see uh you want to take it umm it's like right beyond one of those jersey barriers you want to go in there behind this building where this taxi is coming out.

This instruction includes four separate descriptions of the position of the street ("end of the bridge", "beyond the barriers", "behind this building", and "where this taxi is"); a description of the relative angle of the turn ("hard right"); and a description of the street itself ("hard to see"), which may also warn the driver to devote extra effort to finding it⁶. It is not clear why the subject mentioned the angle of the turn, since there is only one right turn at that place, unless it was either to warn the driver to go extra slowly (for the sharp turn) or because it is hard to see.

One commonly used landmark is the name of the street. Street names make up about one quarter of all landmarks. A street name is not a good clue for when to turn, because signs are hard to see, even in the day light, even when they are present and pointing in the right direction. Nevertheless, drivers do use them. A possible alternative reason for providing a name is to help the driver learn about the city.

Other, equally commonly used landmarks are traffic lights and stop signs.

⁶Note also the use of "very" and "right" to mean "immediate".

EXHIBIT

2

(Part 4)

These are especially useful because the driver is actively looking for them anyway, simply from a desire to avoid accidents.

Notably absent are "famous" landmarks. It is possible to name directions with reference to features such as those named by Kevin Lynch[50] in his classic book *The Image of the City*. These features include widely visible landmarks, nodes (major concentrations of flow or activity), and districts (visually distinct areas within the city). Some subjects did name prominent cultural and geographical features along the route, but only as background, not to tell me when to turn. The landmarks that people did use are what Lynch calls "local landmarks", having meaning only in an immediate context:

2.2.2 Advance Notice

Instructions for an act can be given more than once. One subject gave instructions twice, first in a general form well in advance of the action ("So we're gonna be following Commonwealth for a while and in maybe a mile it bears off to the left and we'll follow it to the left."), and then again at the time of execution ("it bears left here").

Advance instructions may refer to the same landmark more than once. The way that people talk about a landmark depends upon its proximity. Subjects indicate the distance to a landmark implicitly by how they refer to it. While approaching an intersection, the landmarks may not be visible (in this case, the traffic lights are around a curve).

There's a set of lights right up here (gonna go) straight through them
and bear to the right

The subject uses an indefinite article since the objects are not visible. (Note also that two instructions are combined into one utterance.) After going straight

through the lights, the next landmarks are now visible in the distance.

Bear right at that gas station and blinking lights

The subject uses a distal deictic ("that"). Later, when close to the lights, he used a proximal deictic ("these"):

at these flashing lights you'll bear right

2.3 Advice

Cooperative subjects give extra information to make following the route faster and safer. Lane advice tells the driver which lane to drive in when driving on a road with more than one lane. Lane advice can be expressed as positive ("get in the left lane" or "stay in the left lane") to prepare for a turn which can only be made from the selected lane, or as a negative ("make sure you don't get caught too firmly in the left lane") to avoid a lane blocked by traffic waiting to turn or by parked cars, or reserved for turns.

Other advice includes warnings about speed traps (clearly a service that the Back Seat Driver should also provide) and about foolish pedestrians (desirable, but difficult to automate). Other advice is possible. For a discussion of advice about traffic conditions, see chapter 7.

2.4 Style and Packaging

The information that navigators supply includes instructions, advice, and orientation. Above these three kinds of information is a level of style:

People gave instructions in several different ways. Most used simple imperative sentences ("Take a right", "Keep going") while a few used future tense, either second person ("You're gonna take a right") or first plural ("We're gonna turn right"). Still others used an indirect style, talking about the driver's "wants" ("You're gonna wanna take a right"). I think this style is more polite, by speaking as if the goal of getting some place belonged to the driver, not the instruction giver, that is, it's not that the instruction giver is giving orders so as to get to her destination, it's that the instruction giver is providing the driver with information that the driver needs. This style was often combined with a description of the road ("The road forks up ahead."), which justified the choice.

Subjects were fairly consistent in style, but I have no idea why they preferred one form to another. I believe that stylistic choices have more to do with interpersonal relations than with the essentials of direction giving, and interpersonal factors are not accessible through the method of investigation employed in this study. It is essential to this method that people complain about what they dislike. People will complain if instructions are wrong, or even unclear, but they might not complain if they dislike the interpersonal message they think they are getting. (Communication on that level is rarely explicit.) Moreover, I do not believe that my subjects took the computer as a "person", so there were no interpersonal factors. In the section on "mistakes" I describe one way that the Back Seat Driver tries to adopt a gentle style with drivers.

2.5. Silence

We should also consider what is *not* said. At most intersections the driver has a choice of directions; yet we do not find subjects giving instructions at every intersection. This can only mean that subjects assumed that there was a unique choice that was obvious, and that this choice was also obvious to the driver. This

follows from H. P. Grice's maxim of QUANTITY[28]: "Do not make your contribution more informative than is required." Cooperative speakers speak only when required. If a subject is silent, the thing to do must be obvious. Unfortunately, the inverse is not true. That a subject spoke is not evidence that the action described was not obvious. Some subjects give explicit instructions at places where there is no choice whatsoever for example, on one-way streets. Speaking might also serve to reassure the driver of the subject's attention and competence. When we find some subjects speaking, where others are silent, we can guess that the act was not wholly obvious, but we cannot be sure.

The obvious thing is usually to stay on the same road. This may be less obvious when staying on the road requires crossing a bigger road, since subjects sometimes spoke in this situation ("Go straight through the lights"). This follows from the usual pattern for cross-town routes: to go from local streets to progressively larger streets (collectors, then arterials), and then back again to small streets when near the destination. Subjects might want to override a presumption of turning onto the first collector encountered.

At some forks, one branch is the obvious next branch. An example is the connection between Memorial Drive, Brookline Street, and the Boston University Bridge. The right branch (going either up or down the Charles River) leads to a rotary from which one may turn onto Brookline Street or the bridge. The left branch leads up and over the rotary. Most (not all) of those whose routes stayed on Memorial Drive passed through this fork without comment. The left branch is the obvious place to go because it has two lanes and is more straight. An example of a fork without an obvious branch is the connection between Memorial Drive and Massachusetts Avenue, near MIT. Here the left branch leads underneath Massachusetts Avenue, and the right branch forces a turn. Even though the left branch is wider than the right, it is not obvious, perhaps because it departs at a steeper angle than the right branch does. Subjects always were explicit about

which branch to take.

2.6 Example Instructions

Here is one of the transcriptions, and a map showing the route. I have adapted the notation system of Gail Jefferson [72]. This notation system is specialized for study of interaction between speakers. Words are spelled the way they sound, not the way they would be written, so, e.g., "going to" is written as "gonna"; when so pronounced. The notation is explained in table 2.2. I have chosen to capitalize proper nouns to make the account easier to read.

symbol	meaning
curly braces	indicate location
-	hesitation or cut-off speech
L:	the driver is speaking
R:	the passenger is speaking
()	an untimed pause
single parenthesis	enclose uncertain words
double parenthesis	vocal style or non-vocal sound
left bracket	marks simultaneous speech
colon	marks extended length syllable

Table 2.2: Notation system for transcriptions

This route goes from the parking garage at 12 Albany Street, Cambridge, to a garage at Glenville Terrace in Allston.

(at 12 Albany)

- 1 R: and at the stop sign take a right on Main
(on Main Street)
- 2 R: and you'll keep going straight
- 3 R: this is Tech Square
- 4 L: mhhmm
- 5 R: MIT AI Lab is to your left and behind you

6 L: hmm

7 R: so keep going straight past Ames Street

8 R: dont hit the pedestrians

9 L: ((in silly voice)) We want to die;

10 R: (keep) going straight

11 R: merge here I believe with Broadway

12 Go straight at the stop sign

13 and then you're going to take a right on Memorial

14 (at the) sign that says Memorial Drive West

{on Memorial}

15 you can see uh that you're driving right along the river and

16 Boston is on the other side on the left

17 L: mhmm

18 L: what part of Cambridge are we in

19 R: we're in Kendall Square () area ()

20 and basically we're passing by the uh passing by MIT/the long way

21 much of it is right on Memorial Drive

22 L: mhmm

23 L: how long will I be on this street

24 R: oh about another () mile mile and a half

25 L: ((unclear))

{approaching Massachusetts Avenue fork}

26 R: you bear to the left here and go under the bridge

27 under the overpass I guess

28 L: uhhuh

((noise of tires on pavement))

29 R: down the ways a bit we're gonna cross over the river to our left

30 umm cross a bridge called the B.U Bridge

31 but you have to bear off to the right and circle around

32 R: you're gonna want to get into your right lane
33 R: and you're gonna bear off () to the right here ()
34 where it says route 2 ()
35 and this is a traffic circle
36 you're gonna go around three quarters of the way
37 and head across the bridge
38 R: stay in the right lane here cause you're gonna take a right
39 L: uh huh
40 R: so we're leaving Cambridge and going into Boston now
41 and uh Boston University is right around here
42 right along the river
43 L: mhmm
44 R: and we'll be driving past () some of BUs buildings
45 R: you're gonna take the first right here on Commonwealth () Ave
46 R: so we're gonna be following Commonwealth () for a while
47 and in maybe a mile it bears off to the left
48 and we'll follow it to the left
(Commonwealth and Babcock)
49 R: might need to get into the left lane for this uh bearing left
50 ()
51 in fact it bears left here
52 ()
53 middle or left lane
54 R: So I guess a good thing to remember for the directions here
55 is you follow the Green line around (when) you bear left
56 R: now we're gonna pass Harvard Ave
57 and its gonna be the next
58 (no) maybe we're gonna turn onto Harvard Ave
59 come to think of it

60 gonna turn right () on Harvard
61 and that's coming up right here?
62 yes it's the first right then after you bear left
{Commonwealth and Harvard 601}
63 R: okay now we're gonna turn right on the first street
64 which I believe is (glenough)
65 and this is our destination ()
66 towards the end of the street
67 L: here's Glenville on the left
[
68 R: Glen-Glenville
69 L: This street here?
70 R: yes
71 L: Glenville Terrace
72 R: Glenville Terrace

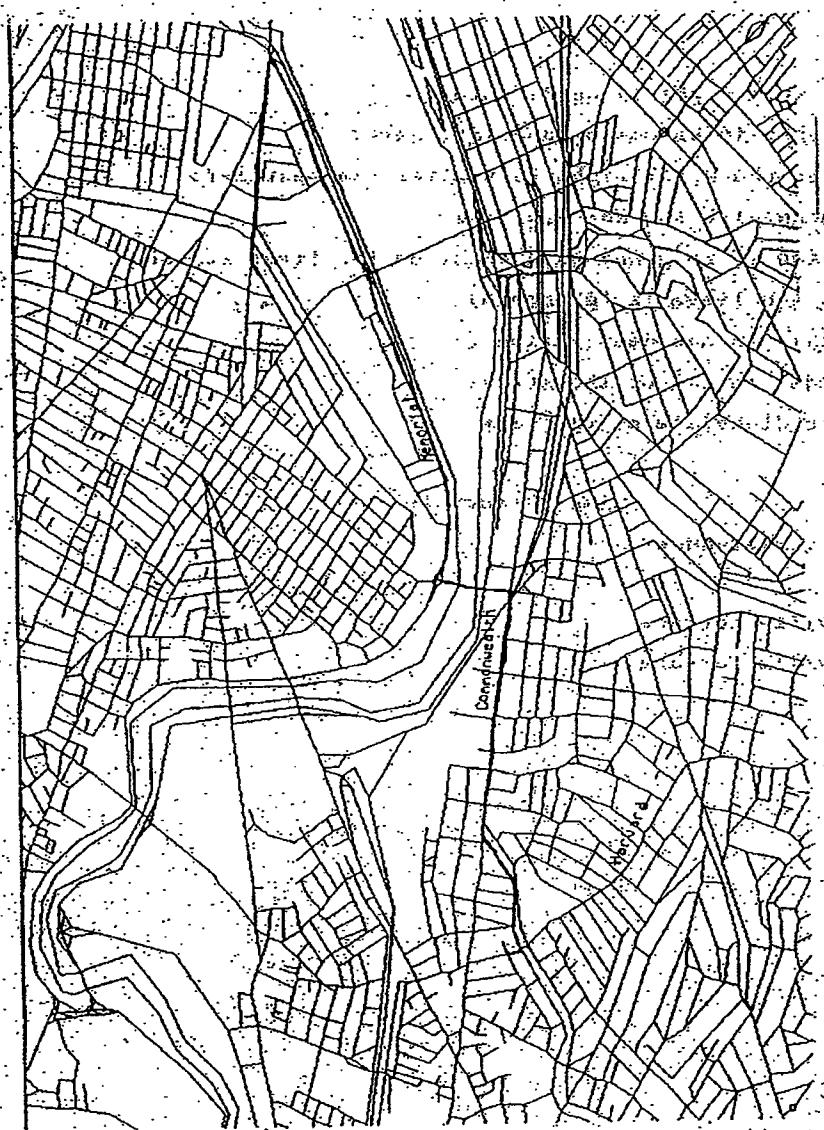


Figure 2-2: Map for sample route

40

0151

Part II

Implementation of the Back Seat Driver

Chapter 3

Cartography

Every navigation system requires a street map. The Back Seat Driver's map originated as a DIME (Dual Independent Map Encoding) file, a map format invented by the US Census Bureau for the 1980 census. Although DIME maps are a useful beginning for navigation, there are problems inherent in the structure which make them unsuited for route finding and route description. One aim of the research here is to determine what information must be in a map database used for these purposes.

The first section of this chapter describes the DIME map format. The second section shows why DIME is not sufficient for route finding or route description. The third section describes the extensions to the map required for the Back Seat Driver, and the final section compares these extensions with those in other navigation systems.

3.1 DIME format

An important point, sometimes overlooked because it seems obvious, is that the

map should be a vector representation, not a raster picture. One reason is economy of storage: the vector representation takes far less room. But the main reason is that only a vector representation is suitable for programmed manipulation. The mathematical basis for route finding is graph theory, not tracing lines on a picture. Even when the intended application is simply to show a picture, the vector representation allows a program to display the map at a variety of scales, detail, and orientation, to highlight significant information and remove unimportant features.

The basic unit of the DIME file is the segment. A segment is a portion of a street (or other linear feature such as a railroad, property line, or shoreline) chosen to be small enough that it is a straight line and has no intersection with any other segment except at its endpoints. The two endpoints are designated FROM and TO. If the segment is a street segment (as opposed to, say, a railroad) and has addresses on it, then the FROM endpoint is the one with the lowest address. Otherwise, the endpoint labels are chosen arbitrarily. A segment has two sides, left and right. The sides are chosen with respect to travel from the FROM endpoint to the TO endpoint.

Attributes of a segment include:

- its name (40 characters)
- its "type" (a one to four character abbreviation such as "ST")
- longitude and latitude of the endpoints
- ZIP code for each side
- addresses for each endpoint and each side
- list of segments connected at each endpoint

A program using DIME can find the location of an address along the segment by interpolating the addresses between the low and high addresses for the two endpoints.

3.2 The limits of DIME

The form and content of a map database depend upon the purpose for which you intend to use it. The DIME file was invented to allow the Census Bureau to determine the proper census tract for any given address in the USA, and for no other purpose. The DIME file is well suited to determining the absolute position of a building from its street address. But the DIME database is not sufficient for route finding, and it is only marginal for generating route descriptions. The problems with DIME fall into two categories:

3.2.1 The model of connection is deficient

The DIME file indicates that two segments are *physically connected* (that is, they touch), but not whether they are *legally connected* (i.e. it is legal to travel from one to the other). Legal connectivity is crucial for route finding, unless we only want to go to jail or the hospital. It must be explicit. Legal connectivity does not replace physical connectivity. Only legal connectivity is needed for route finding, but route description requires information about physical connections as well. Physical connectivity also affects route finding directly when seeking the simplest route, since ease of description is determined in part by physical connectivity.

The DIME file is a planar graph. This means that no two segments can cross except at an intersection, so there is no way to correctly represent, say, an overpass. The DIME format represents an overpass by "breaking" both streets at the point where they cross, and creating a fictitious intersection even though the segments do not touch in reality. These false intersections are particularly troublesome since DIME does not represent legal connectivity, so it appears possible and legal for a car to jump straight up and turn onto the overpass.

3.2.2 Position resolution is inadequate

Coordinates in DIME are stored in ten thousandths of a degree. This means that the position of an endpoint in the map differs from the true position by as much as 6.5 meters in latitude (north/south) and 5 meters in longitude at the latitude of Boston. (The size of a degree of longitude depends upon latitude, since longitude lines converge at the poles.) This inherent position error causes problems because it introduces error in length and in heading. See figure 3-1 for an illustration. Here, the points A and B could have been located anywhere within

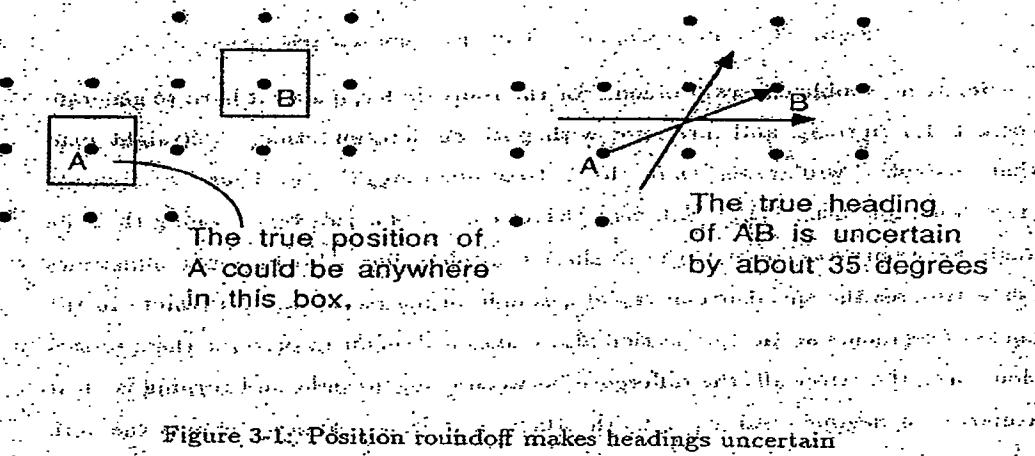


Figure 3-1. Position roundoff makes headings uncertain

the boxes surrounding them. Depending where A and B really are, the heading of line AB changes by almost 90 degrees. Figure 3-2 shows how the accuracy in heading decreases as segments become shorter.

Uncertainty in heading causes uncertainty in the angle between two segments. A straight street can appear to wobble if it is made of many short segments. See for instance Pleasant Street in figure 3-3. The dots in this figure are the possible coordinate locations. Pleasant Street lies in between two columns; so it wobbles back and forth.

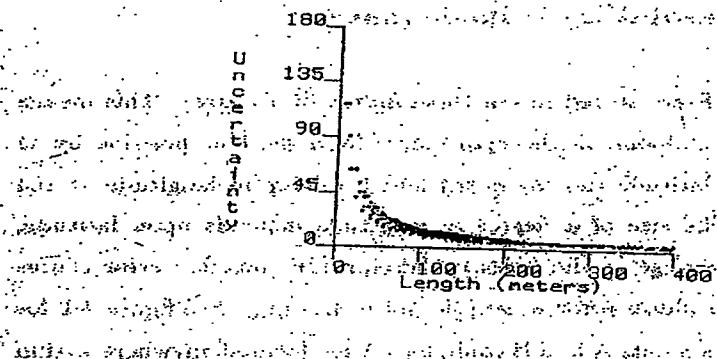


Figure 3-2: Short segments have less angular resolution.

Segment wobble causes problems for the route finder, makes it hard to generate correct descriptions, and interferes with position determination. A straight road that "wobbles" will appear to be slower than one that does not because cars must slow down for turns. The Back Seat Driver corrects for this by assuming that the angle between two streets is the smallest possible value. This means it sometimes overestimates the speed it can travel through an intersection. Uncertainty in the angle of segments at an intersection also makes it difficult to describe the intersection correctly. After all, the difference between going straight and turning is "just a matter of degree". As a result, the Back Seat Driver sometimes uses the verb "bear" where "turn" would be more appropriate, because it underestimates the angle of the intersection. Segment wobble also interferes with navigation because it makes it difficult to compare compass headings with the heading of a street.

3.2.3 The forthcoming TIGER format is better, but not enough

The Census Bureau, in cooperation with the United States Geological Survey,

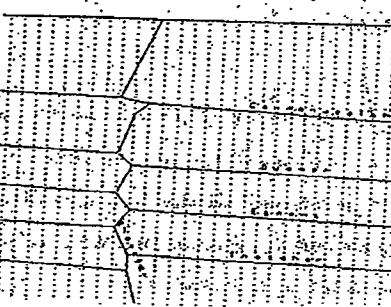


Figure 3-3: Coordinate rounding makes streets seem to wobble.

has designed a new map format known as TIGER (Topologically Integrated Geographic Encoding and Referencing). This format has several improvements from the DIME format[79, 52; 42]. The map format is structured to permit automated consistency checking, so it should be more reliable[13, 54]. TIGER coordinates have two more decimal places of precision (bringing accuracy to something like four inches), and the segment representation includes "curve vectors" so that a segment need no longer be a straight line. Two dimensional areas are explicitly represented as closed polygons, which will make it much easier to tell whether a route crosses a river. The database includes point landmarks (fire towers, churches, schools) and "Key Geographic Locations" (mostly commercial buildings)[53] and this will make it easier to find locations and give directions. But the TIGER file is still a planar graph and has only physical connections, so it will not be sufficient for route finding. TIGER maps will be an important component of future navigation programs because they are more up to date and more accurate, but will still be insufficient. Those who wish to use TIGER files will have to obtain and represent connectivity information in some other way.

3.3 A better map

Implementing the Back Seat Driver required me to extend the map format. Each extension described here is used in one or more ways by the Back Seat Driver. Since I have not yet described all the features of the Back Seat Driver in detail, it may not be clear why certain extensions were required. I ask the reader's patience if certain decisions seem unmotivated.

3.3.1 Legal connectivity

The most significant addition is an explicit representation of legal connectivity. This is crucial for finding routes. At each endpoint of the segment is a list of all segments which are legally accessible from that endpoint. This list allows the route finder to consider only legal paths. Although this change is easy to implement, it does not seem to have been included in any other navigation system.

3.4 Additional segment attributes

The DIME file records a small amount of information about the segment. I found it useful to add additional attributes to the segment to make better descriptions. These new attributes are:

- street quality
- divided roads
- signs
- other landmarks
- traffic lights

- stop signs
- lane information
- speed limit

The first four were added for Direction Assistance, the remainder for the Back Seat Driver.

3.4.1 Quality combines size and speed of a segment

The street quality is a number from 1 ("super") to 4 ("bad") which combines the ease of locating and following the street and the expected rate of travel along it. Super streets are the access highways (e.g., the Massachusetts Turnpike, and most of Storrow Drive, but not Memorial Drive). Figure 3-4 shows the network of super streets. In general, the "super" streets are well known, though perhaps not by their correct name. The "bad" streets may also be well known, to local residents, but not by name. (For instance, see "Back Street" on page 31.) The street quality attribute affects the route finder and the route describer. The route finder ordinarily prefers the street of highest quality. The route describer uses the quality of a street in several ways, described below.

3.4.2 Expanded street classifications

It is useful to have a richer taxonomy of street types than that provided by DIME. There are five categories of streets: ordinary, rotary, access ramp, underpass, tunnel, and bridge. There are also four categories of non-streets: railroad, water, alley, and walkway.

A rotary (or roundabout, or traffic circle) is a kind of intersection in the form of a one way circular road (figure 3-5). Traffic enters and exits the rotary on lines

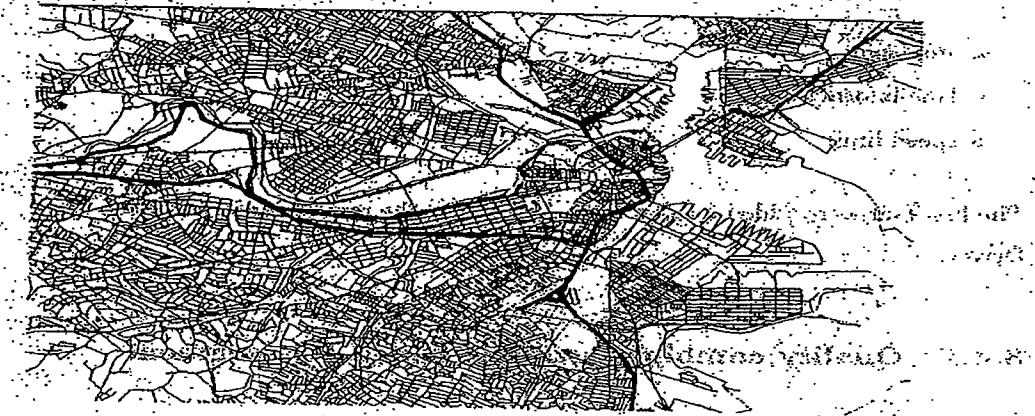
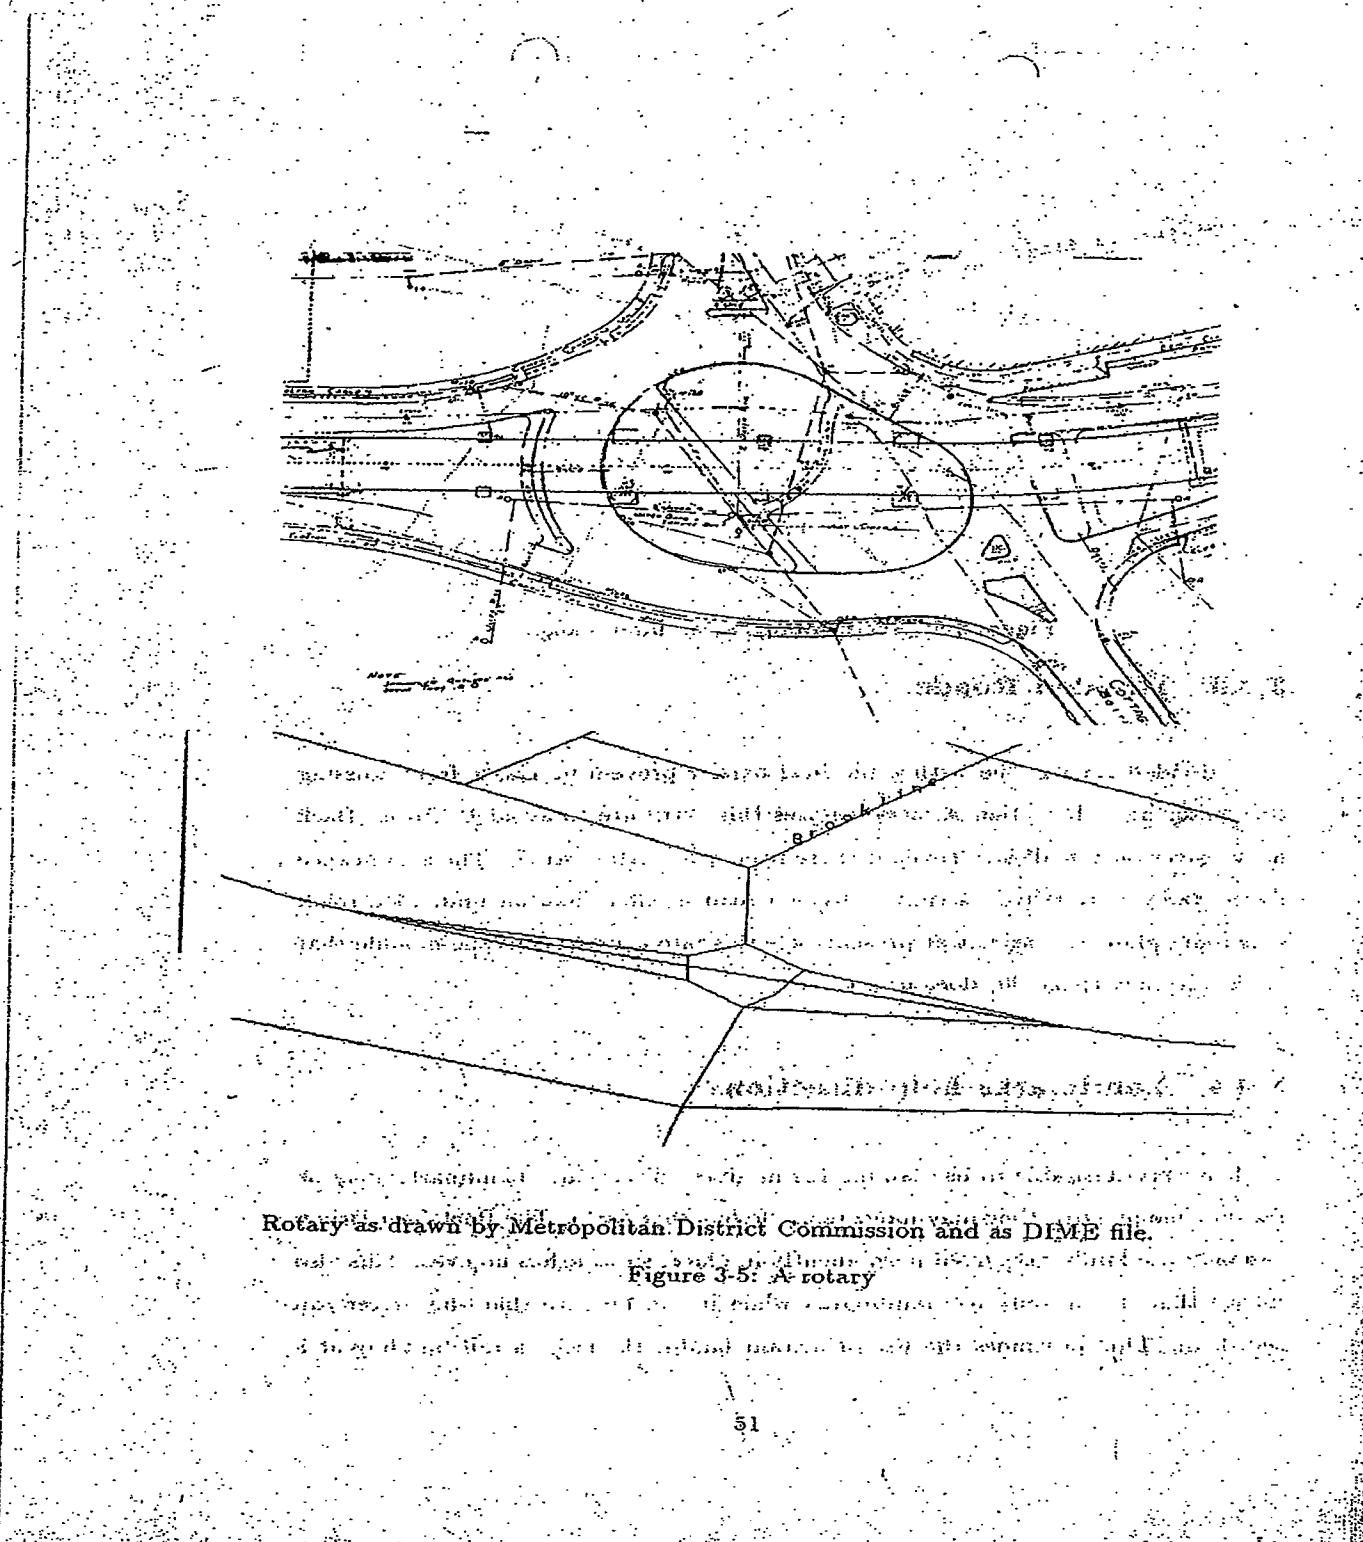


Figure 3-4: The "super" streets of Boston

more or less tangent to the circle. Rotaries are efficient, since traffic can move continuously instead of alternating, and safer than ordinary intersections, since all traffic is moving the same direction, with a low relative velocity. On the other hand, rotaries require more space, and can be somewhat confusing. The Back Seat Driver needs an explicit indication of whether a segment is part of a rotary, since it would be too difficult to guess from the map geometry.

An access ramp is a nameless piece of road that leads to or away from another street. The prototypical access ramp is the highway entrance or exit, but access ramps can also link interchanges between ordinary streets. Access ramps are almost always one-way streets, and they have no address numbers. It is important to distinguish access ramps because they must be described differently than ordinary streets.

The classes underpass, tunnel, and bridge have obvious meanings. They are helpful in route descriptions because the driver is very likely to notice these features, and in forming names, since bridges and tunnels have singular names (*the Harvard Bridge, the Sumner Tunnel*).



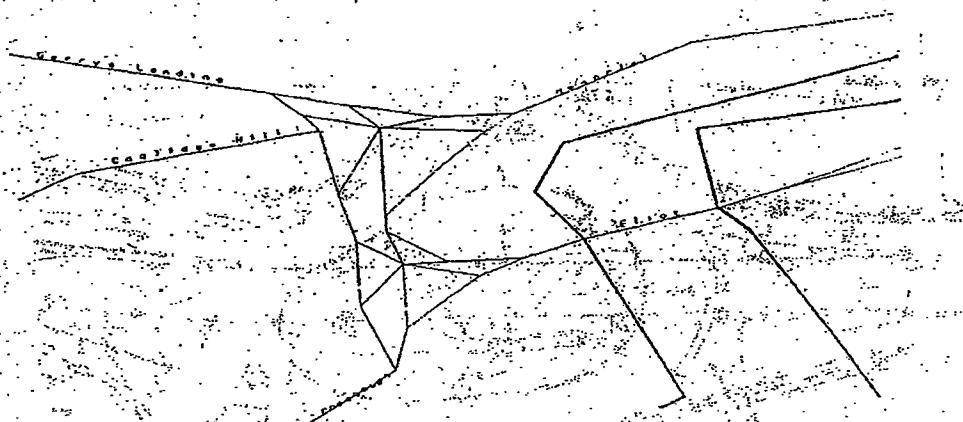


Figure 3-6: Access ramps at an interchange

3.4.3 Divided Roads

A divided road is one with a physical barrier preventing traffic from crossing the center line. Direction Assistance uses this attribute to avoid U Turns (Back Seat Driver only makes U Turns if there is no other alternative). The significance to the Back Seat Driver is that a divided road is safer than an undivided road. The route planner does not at present take this into account, but the module that gives warnings (page 90) does use it.

3.4.4 Landmarks help directions

It is very desirable to use landmarks in giving directions. Landmarks may be fixed in one place or they may move (e.g., "follow that car"). The Back Seat Driver can only use landmarks fixed permanently in place, since it has no eyes. This also means that it can only use landmarks when it can be sure that the driver can see them. This precludes the use of distant landmarks (e.g., a tall building at a

distance) because it can not tell whether such a landmark is visible at any given place.

The most useful landmarks are traffic lights. There are two kinds of lights, three way ("stop") lights and blinking lights. Traffic lights are stored independently for each endpoint of each segment, since the presence of a light at one segment of an intersection does not imply that all others do.

Signs are another class of landmarks. These are especially useful on highways, where exits have both a number and a name. Highway signs are designed to be visible under all conditions and at great distances, so they are reliable landmarks. A sign or exit number is stored as a connection cue, which is a text string that gives a cue for moving from one segment to another. Every cue has a type, which tells the kind of cue, e.g. sign or exit-number. There may be more than one connection cue for a given pair of segments, but there is never more than one of a type. A problem with this representation is that the information in the sign is stored as unstructured text. It is important that the Back Seat Driver understand what the sign says, not simply utter the words. There are two reasons for this. First, the Back Seat Driver's internal representation for text is based on syntactic structure, not text strings. Second, the objects mentioned in the signs (cities and roads) should be entered into the discourse model. They should become salient for future reference. This means that the text of a sign must be parsed, so that e.g. the sign text "Cambridge, Somerville, and Storrow Drive" should become a conjunction of the two cities "Cambridge" and "Somerville" and the street named "Storrow Drive". The Back Seat Driver parses sign text by separating it into tokens delimited by commas and the word "and", then attempts to recognize objects on the map (street names, cities, neighborhoods) from these tokens. When recognition fails, the token is not entered into the discourse model. When parsing fails, the spoken output has incorrect grammar.

Buildings are a third class of landmarks. The two types of buildings used at present are toll booths and gas stations. Toll booths and gas stations are stored in two different ways. Toll booths were added to the map database for Direction Assistance and are stored as connection cues. This is a bad representation for buildings for the same reasons that it is bad for signs, but is used for historical reasons. Gas stations are stored in the Yellow Pages file (page 88). This is better, but costly, as it requires searching the file for gas stations. A better approach would be to index gas stations (and other buildings) by street, as is used in TIGER (page 46).

3.4.5 Lane information

Roads often have more than one lane. Selecting the proper lane can make travel faster, and it may even be mandatory, since certain turns may only be possible from some lanes. For both directions on the segment is recorded the number of lanes and whether one or more lanes is reserved for turn restrictions - either left turn only or right turn only.

3.4.6 DIME files need correction

As distributed by the Census Bureau, the DIME file for Boston has many errors. Errors arise from imprecise surveying, inconsistent coding, clerical errors in data entry, and changes since the survey date. The map is especially bad for access ramps, which often have gross errors in angle, or are missing altogether. This is probably because the map makers had little incentive to be accurate with roads where no one lived, because there are no census data to be taken there.

Names are often inconsistent, either because of spelling errors or because a portion of the name is omitted (e.g. "North" in "North Beacon Street".) To use

the Boston DIME file required reexamining the entire area in person. For the most part, this required no more than an hour for each square mile. More time was required for complicated intersections which had the kind of anonymous roads so poorly surveyed in the DIME maps. To enter these, I would make a sketch map on the spot, then enter the changes using a graphical editor. This allowed for reasonably quick data entry, but was not very accurate, since I used only my eyes to measure distance and angle.

I had edited much of the map earlier for Direction Assistance, but I also found that the Back Seat Driver requires more accuracy than Direction Assistance because the Back Seat Driver needs accurate positions to deliver messages at the right time. Also, people following written directions (as in Direction Assistance) rely more on their own intelligence to figure out when to act. If there is a discrepancy between the instructions and the world, they look around to try to understand it and correct it. But users of the Back Seat Driver just do what the machine tells them to do.

3.5 Other maps

Those who have built other automotive navigation systems have also found it necessary to add features to their maps. Here I describe features beyond those present in DIME.

Neukircher [58] describes the features of the map used in the EVA system. This map has better position information than DIME. Points are stored in three dimensions and are accurate to 2.5 meters. Road segments are straight lines, chosen so that a new segment begins at either an intersection or when the change in direction exceeds 30 degrees, or when the distance from the center line exceeds 5 meters. Additional attributes of the roads include height and weight restrictions.

location of magnetic anomalies, warnings, landmarks, special objects useful in descriptions (e.g., underpasses), layout of complex intersections, and signs. The map has two levels of detail[65]. The coarse level is used for route finding, and the fine level has more detailed information for position finding. Route finding information includes *two* values for expected speed (one for normal conditions and a second for times of high density), the expected wait time at segment endpoints, and areas where children are likely to be playing.

The University of Calgary AVL-2000 system uses a map that originated as a Canadian government Area Master File. This format, similar to DIME, also required extensive augmentation[33]. Link (segment) attributes include distance, expected travel time, safety, scenic value, tolls, "impediment value" [sic], one way limitations, banned turns, road type (over- and under-pass, traffic circles, clover leaf), presence of meridians (divided?), and restricted areas. Harris describes as a "special problem" those "source or destination points which correspond to a street addresses [sic] which do not have a unique node identifier". Either their map representation cannot interpolate addresses along segments, or the route finder is restricted to finding routes to nodes only. Harris also mentions auxiliary road information including landmarks, points of interest, emergency services, commercial establishments, weather conditions, traffic flow, and road characteristics, and stresses the importance of being able to update the map database over a communications link while driving.

Most systems have expanded the classification of streets. The ETAK map classification is:

- Interstate highway
- Semi-limited Access Roads and State Highways
- Arterial
- Collector

- Light Duty Roads

- Alleys or Unpaved Roads

- High Speed Ramps

- Low Speed Ramps

This is a richer taxonomy than that of the Back Seat Driver, and is essential to EТАIK because of its significance in choosing which roads to display (lesser roads are suppressed at larger scales to control detail) and in which colors to display them. The Back Seat Driver could also benefit from such a taxonomy. The EVA system has a two level taxonomy:

- rural
 - motorways and federal highways with separate directional lanes and without intersections
 - federal highways
 - roads wider than six meters
 - roads four to six meters wide
 - others
- urban
 - divided
 - through
 - main
 - side
 - restricted

These maps have some questionable design decisions on the representation of legal restrictions. The EТАIK map has no legal topology at all. It is not intended for route finding. The EVA map apparently encodes restrictions on turning by signs, rather than directly in the network. The Calgary map represents legal topology (one ways, banned turns) as a link attribute instead of in the network.

topology. It may be that the street network represents one physical topology, with the assumption that legal topology will be equivalent to physical topology unless specially indicated. The Back Seat Driver appears to be unique in maintaining separate but equal representations for physical and legal topology. These two topologies should be integrated because legal topology is needed for route finding, and physical topology for route description.

Some other navigation systems attempt to give warnings about hazardous conditions. In those systems (EVA, Calgary) the hazards (about slope, width, or curves) are encoded explicitly into the map. It is unclear how "safety" (in the Calgary map) is represented, or whether it *should* be explicitly represented. A system which can compute safety from other attributes (traffic flow, slope, width) will be more general than one which relies on stored data. On the other hand, if safety is based on external information (accident statistics) then it belongs in the database.

3.6 Improvements are still needed

The Back Seat Driver's map, advanced as it is, could still be extended. Both route finder and descriptions would benefit from a more powerful map database. Desirable enhancements include:

- Time dependent legal connectivity. Sometimes a given turn will be prohibited only at certain hours of the day, typically rush hour. At present, Back Seat Driver must record such turns as always prohibited. This results in less than optimal routes. In Cambridge, Memorial Drive is closed on Sundays during the summer to create a large play space. The Back Seat Driver map representation has no way to record such arbitrary, yet predictable changes.
- In general, the Back Seat Driver needs the ability to change legal connectivity.

dynamically, since a street can be closed suddenly because of an accident or weather. This is more than just a database issue, so I discuss it in Chapter 8.

- **Expected rate of travel.** Rate of travel is at present taken to be a function of street "quality". This is probably a mistake. Although there is a correlation, travel rates should be a separate, explicit segment attribute. One reason is that travel rate, unlike quality, changes during the day. For example, the Central Artery in Boston is a "super" street, but is also very slow during rush hour. Either the Back Seat Driver should have a model of traffic flow, like that of an experienced driver (i.e. it should know what "rush hour" means), or it should have some means of getting real time traffic conditions, perhaps broadcast by a central reporting agency. I discuss this further in Chapter 8, since I see it as more of a planning issue than a database issue.
- **Turn resistance.** Some turns, though legal, are also difficult to make (e.g. turning left across a large traffic flow at an intersection without a traffic light). The route finder should avoid these turns if possible. To an extent, the difficulty of a turn is implicit in the quality of the participating street segments, but an explicit model might prove useful. The TAXI! driving simulator[12] has the concept of turn resistance, which is a number from one to ten specifying the difficulty of making a given turn. This concept should be adopted.
- **Lane restrictions.** Some lanes or streets are restricted to certain kinds of traffic (car pools, no commercial vehicles). Including this information would make the Back Seat Driver useful for more kinds of traffic. An important subclass is height restrictions, since some underpasses are so low that tall vehicles will not fit under them. (The Memorial Drive underpass under Massachusetts Avenue is an infamous example. See figure 3-7.) The route finder should not attempt to send tall vehicles to their doom!

¹I assume that the route finder knows about the height, weight, etc. of the vehicle in which it is installed.

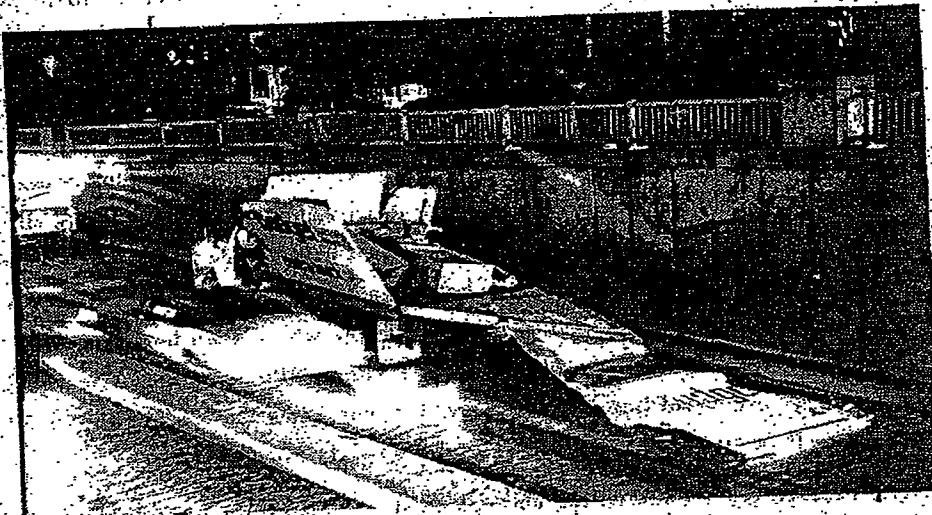


Figure 3-7: When trucks ignore height restrictions

- **Traffic light information.** At some lights it is permitted to make a right turn at a red light after a full stop. Right turns here will be no slower than right turns at a stop sign, so the route finder should prefer such intersections to those that do not permit it. Also, traffic lights have differing cycle lengths.
- **Altitude.** Points on street maps should be three dimensional. Route descriptions would be better given knowledge of the underlying topography ("Go up the hill, and take a left"). Stopping distance is affected by slope, so instructions must be given sooner when traveling down a hill. Slope affects safety. The route finder should avoid steep slopes in snowy weather. Knowledge of altitude provides a constraint for GPS position determination. Given exact knowledge of altitude, three satellites suffice for determining horizontal position. Less exact knowledge might still be useful. Finally, distance

EXHIBIT

2

(Part 5)

between points will be more accurate when change in altitude is considered.

- At grade or not. Roads designed for high speed may be more level than the underlying topography. They may be elevated or they may be depressed. A road which is not at grade will not have the slope of the land beneath it.
- Local knowledge is a catch-all term for facts about the road which can not be predicted from the map. These are facts about how people and institutions acts on or near the road; e.g. that a speed trap is here, or that this road is one of the first ones plowed after a snow storm.

3.6.1 Summary

A street map used for route finding must include a representation for legal connectivity. This is essential. To find the fastest routes, the map should also include features that affect speed of travel, including street quality, speed limit, traffic lights and stop signs. To generate directions, the map should include landmarks such as traffic lights and buildings, and additional descriptive information about the street segments, including street type, number of lanes, turn restrictions, street quality, and speed limit. The Back Seat Driver map includes all of these.

It would also be desirable to include other features, such as time dependent legal connectivity, and expected rate of travel along streets and across intersections. Positions should be stored in three dimensions, not two, and with sufficient accuracy that the headings of segments can be accurately determined from the map. These features should be added in future work.

Chapter 4

Descriptions

The two key issues in describing a route are deciding what to say and deciding when to say it. There is a tradeoff between these two factors. At one extreme are directions given completely in advance, with no control over when the driver reads them. A directions of this kind might be "Go half a mile, then take a left onto Mulberry Street". A driver following such an instruction must use the odometer to estimate distance or look for a street sign. The instruction itself does not say when to act. On the other extreme are instructions which rely totally on timing for success. Such an instruction might be: "Turn left now."

The more the Back Seat Driver uses timing, the less the workload on the driver. Where possible, instructions should use timing as much as possible, and contain only as much description as necessary. The limiting factor is the accuracy of the information the Back Seat Driver has about position. For instance, consider a turn to the right where there are two roads to the right, one a sharper turn. If the Back Seat Driver has sufficient directional resolution to tell which way the car is facing (and if the map is correct) it can designate the correct road as the driver is turning. If the driver is headed the wrong way it can say "no, the other one".

The current Back Seat Driver does not have this kind of information, so it must describe the turn, saying either "a sharp right" or "an easy right".

The Back Seat Driver tries to minimize the driver's workload. It does not rely on the driver to measure distances. This is just as well, since many drivers do not even know what an odometer is. More important, the driver has a lot more to do than keep track of distance and heading - namely keeping the car on the road and avoiding collisions, not to mention listening to the radio, talking to someone, or just thinking. Minimizing the driver's workload means telling the driver when to act, not just what to do.

The next section tells how the Back Seat Driver forms its descriptions. The basic process is to look ahead at the route and decide what action the driver must take to follow it, then form sentences that tell the driver about these actions, uttering them at the appropriate times. Following sections take up some additional issues, namely giving appropriate advice while following the route, and handling mistakes when the driver does not follow the instructions.

4.1 Classifying movement

To the Back Seat Driver, a route is a sequence of street segments leading from the origin to the destination. Each connection from one segment to another is considered an "intersection", even if there is only one next segment at the intersection. At any moment, the car will be on one of the segments of the route, approaching an intersection. The task of the Back Seat Driver is to say whatever is necessary to get the driver to go from the current segment, across the intersection, to the next segment of the route. Most often, nothing need be said. But at other times, the Back Seat Driver will need to give an instruction.

Instructions must use terms familiar to the driver. An example is what to say at a fork in the road: Considering only topology, there is no difference between a fork and a turn, but it would be confusing to call a fork a turn.

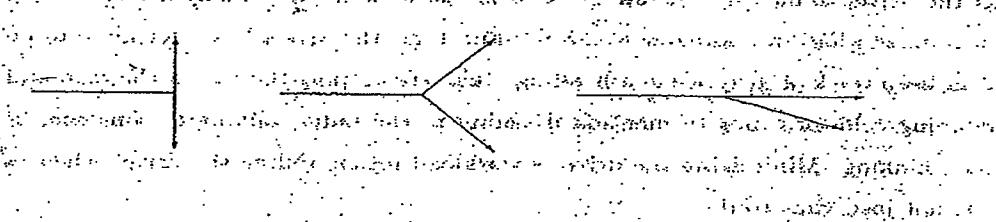


Figure 4.1: T, fork, and exit all have same topology
The difference is geometry, i.e. the angle between the two segments. Likewise, the difference between a fork and a highway exit is not a matter of geometry, but requires knowing that the exit is an access ramp.

The Back Seat Driver has a taxonomy of ten intersection types. An intersection type is called an act because the important thing about an intersection is what action the driver takes to get across it. The Back Seat Driver is implemented with an object oriented programming methodology, so for each act there is an "expert" capable of recognizing and describing the act.

The Back Seat Driver generates speech by consulting these experts. At any moment, there will be exactly one expert in charge of telling the driver what to do. To select this expert, the Back Seat Driver asks each expert in turn to decide whether it applies to the intersection. The experts are consulted in a fixed order, the most specific ones first. The first expert to claim responsibility is selected. This expert then has the responsibility of deciding what (if anything) to say.

Each act has a recognition predicate. A predicate can consider topology, geometry, the types of street involved, or any other factor. The predicate also

decides whether the move is *obvious*, that is, the driver can be trusted to do it without being explicitly told to do so. Actions that are obvious are not described.

If the next action is obvious, the Back Seat Driver looks ahead along the route until it finds one which is *not* obvious. There will always be at least one, because stopping at the end is never obvious.

The actions are:

- **CONTINUE**
- **FORCED-TURN**
- **TURN-AROUND**
- **ENTER**
- **EXIT**
- **ONTO-ROTARY**
- **EXIT-ROTARY**
- **WORK**
- **TURN**
- **STOP**

4.2 Definitions of actions

A **CONTINUE** is a place where the driver stays on the "same" road. Almost always, a continue is obvious. This act is the first on the list to catch to cases where nothing should be said. The continuation of a street depends on the type of street:

- from a road, it is the next road segment.
- from a rotary, it is the next rotary segment.

- from an access ramp, if there is exactly one next segment, that is the continuation; otherwise there is no obvious next segment.
- otherwise, it is the one segment that requires no more than 30 degrees of angle change (if there is exactly one, and if it is not a rotary) or the one segment with the same name (if there is exactly one). The reason for comparing names is *not* because the driver is aware of the name, but because the designer who named the street was. The assumption is that if two segments have the same name, they are the same street, and that is why they have the same name. This "sameness" is presumably reflected in details not captured by the map, for example continuity of painted centerline. There are many places in the area where the obvious "straight" continuation of a segment is at an angle as great as 45 degrees, but it would not be right to call this a turn.

In almost all cases, to continue is the "obvious" thing to do. However, some drivers asked for explicit instructions before continuing across an intersection with another, larger street. This makes sense, since a common pattern for routes is to start on local roads, move onto larger roads (collectors and arterials), and then back again to local roads near the end. Drivers who ask for confirmation at major intersections are noticing that they have come to a major choice point.

In general, it would be wrong to expect to find a single universal definition for obviousness. The Back Seat Driver has a limited ability to be adjusted to suit the idiosyncrasies of the driver. This ability is in the form of a simple user model, which is a set of parameters about a particular user. For instance, one parameter specifies whether the user needs confirmation when crossing a major intersection. The user model can only be adjusted by a programmer, not by the user, and not by the system itself, either. A major topic for further work is to learn how a system like the Back Seat Driver can learn for itself the user's preferences and desires. Despite its simplicity, the user model makes the Back Seat Driver more comfortable for its drivers.

A FORCED TURN is an intersection where there is only one next street segment where the road bends more than 10 degrees. Even though there is no decision to make at a forced turn, it is useful to mention because it strengthens the driver's sense that the Back Seat Driver really knows about the road conditions. (It may be helpful to think of "obvious" as meaning "worth mentioning". It is worth mentioning a bend in the road, even though it is also obvious that one stays on the road.) The natural directions collected at the start of this study also included examples of "forced turns". A forced turn is not worth mentioning if both segments are part of a bridge, a tunnel, or an access ramp, or if the angle is less than 20 degrees. The intent here is to estimate whether the driver can see the continuation.

The TURN AROUND action is recognized when the heading of the car is the opposite of what it should be. Recall that a route is a sequence of segments and endpoints. At all times the car will be on one of the segments in the sequence. If the car's orientation is not the same as the endpoint in the path, then the driver must turn around. Note that the route finder only calls for a U Turn if there is no other way (e.g. when facing into a dead-end street).

The next four actions depend heavily on the street type and street quality in order to be recognized correctly. This must be emphasized because these features are not always present in digital street maps, yet without them, these acts can not be identified. (One could imagine attempting to infer the presence of a rotary from the geometry of streets in a map. It seems very difficult.)

To ENTER is to move onto a super street (or an access ramp that leads eventually to a super street) from an ordinary street, but not from a super street or an earlier access ramp. Similarly, to EXIT is to move from a "super" quality street onto a street with lesser quality that is either an access ramp, or has a different name. The extra condition is needed because some "super" roads are not uniformly "super": for instance, the McGrath Highway in Somerville is a limited access road in places, but has stoplights at other places. It would not be right to call the

change in quality an "exit". The name criterion does not apply to access ramps because their names have no significance.

To go ONTO a ROTARY is just to move from a segment that is not a rotary onto a segment that is, and to EXIT a ROTARY is to do the reverse. Again, this is an act which can be correctly described only if the street map database includes an explicit marking of streets as rotaries.

At a FORK, there must be at least two alternatives, all within a narrow angle, and none of the branches must be the obvious next segment - that is, the branches must all be more or less equal. Either all the alternatives must be access ramps, or none of them must be. One branch is "obvious" if it is the only branch with the same level of quality, or if it is markedly straighter than the others, or if it is the only one with the same number of lanes; provided that all of these clues agree. If one branch is stronger than the others, the intersection is not a FORK. It is either a CONTINUE or a TURN.

The STOP action is recognized when the car is on the destination segment.

Finally, a TURN is anything not handled by one of the above cases.

The greatest weakness of this approach is that the classification predicates are sensitive to small changes in the angles between segments. It is not likely that people use absolute numbers (e.g. 10 degrees) as cut-off values for their determination of how to describe an intersection. More likely, different classifications compete.

Still more important, people making classifications will use visual cues, not just facts from the map.

By classifying an act, the Back Seat Driver selects an "expert" to describe it. The driver needs to know what to do and when to do it. The next section tells how the experts decide what to say to convey this information. A general concern, not limited to any particular area, is to give the driver the impression that the program is actually present in the car, and sees the road in the same way the driver does. As will be seen, this concern has consequences in several different places.

4.3 What to do

Each action has a description function to generate a description of the action. The description function takes inputs specifying the size of the description (brief or long), the tense (past, present, or future), and the reference position. A short description is the minimum necessary for the act. It is typically an imperative (e.g. "Bear left"). A long description includes other facts about the action, an expression indicating the distance or time until the act is to be performed, and possibly information about the next act, if it is close. The reference position is a position (along the route) from which the action is to be described.

To motivate this discussion, here's a sample of the description of the left turn from Fulkerson Street to Main Street in Kendall Square, Cambridge (as seen in figure 4-2):

Get in the left lane because you're going to take a left at the next set of lights. It's a complicated intersection because there are two streets on the left. You want the sharper of the two. It's also the better of them. After the turn, get into the right lane.

This instruction begins with a piece of lane advice: an action to be taken immediately, then describes an action in the near future. The action is a TURN, though that word is not used explicitly. It tells the direction of the turn (left) and specifies a landmark (the lights) that says where the turn is. In many cases, this would be enough, but here there are two streets on the left, so the instruction goes on to specify the desired road in two ways (by comparative position and relative quality). Finally, it concludes with some lane advice to be executed during (or just after) the act.

This is the most complicated text that the Back Seat Driver has produced. Remember length and detail are *not* virtues in giving directions! The Back Seat

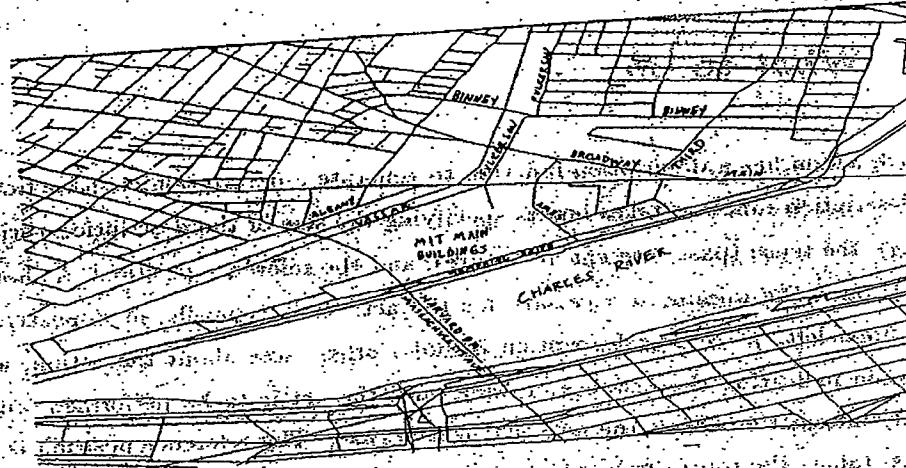


Figure 4-2: Map near MIT

Driver produces this long a text only because it has no better means to make the driver follow the route. If a shorter text would accomplish the same aim, it would be better.

A brief description consists only of a verb phrase. The verb depends on the type of act and perhaps on the specifics of the act. For instance, if the turn described above were at a less sharp angle, the verb phrase would have been "bear right" rather than "take a right". Besides the verb itself, the verb phrase must say which way to go. In most cases, the word "left" or "right" is sufficient. Where it is not, the possibilities are to use a landmark or to describe the turn. A landmark can be either in the appropriate direction ("towards the underpass") or the other direction ("away from the river"). Specifying direction with a landmark has the advantage that some drivers confuse left and right sides, or mishear the words, so it is a redundant cue. Also, it increases the driver's confidence that the system really knows what the land looks like. A description of the turn can mention either quality or the relative angle of the desired road. The angle must be described qualitatively (more or less "sharp"). It would be more precise to use the angular distance (e.g., "turn right 83 degrees"), but drivers would not understand it.

Besides telling drivers what to do, the Back Seat Driver must also tell them when to do it. The Back Seat Driver has two ways to do this. If uses timing ("Take a left here") when the driver has reached the place to act. When the act is more than a few seconds in the future, it uses a long description, which includes one or more cues which either describe the place for the act, the features of the road between the current location and the place, or the distance of time until the act. The next section describes these.

4.4 When to do it

Good directions tell the driver where or when the next action will be. One way to do this is by speaking at the moment to act, and this is the subject I discuss first. But it is also useful to give instructions before the act, if time permits. This allows time for preparation, if required, permits the driver to hear the instruction twice, and also spares the driver the need to be constantly alert for a command which must be obeyed at once.

4.4.1 Timing

The system gives instructions at the time the act is to be done. This is the most important time to give instructions. The Back Seat Driver does not assume that the driver will recognize the place to act (e.g., by seeing a street sign) so the driver must be told when (or where) to act. The system calculates the place to begin speaking by finding a distance back from the intersection which is $v * (t_{speak} + t_{reaction})$, where t_{speak} is the time to speak the utterance and $t_{reaction}$ is the driver's reaction time. The time to speak depends on the number of words in the utterance. (The synthesizer speaks 180 words per minute.) Reaction time is taken to be two seconds.

The system speaks autonomously, but can also speak on demand. At any time the driver can push one of two buttons ("what next?" and "what now?") to ask for instructions immediately. The first button I added was "what next?" which gives the next instruction immediately. In testing the Back Seat Driver, I found that drivers hit this button when they were unsure of what to do at the time they hit it, either because they had come a place where they did not know what to do, even though the system thought it was obvious, or they were unsure whether they had come to the place where the system expected them to act. What they really needed was a "what now?" button. The "what now?" button gives the description for the very next segment transition, regardless of whether it is obvious. This function is placed on the "1" button on the cellular phone keypad, so it can be found by touch while driving.

4.4.2 Cues

A cue gives either a description of the place for the next action, or the distance or time to it. This description should be so clear that the driver cannot only recognize the place when it comes, but can also be *confident* in advance that she will be able to recognize the place. They must not only be clear; they must also seem clear[68]. Appelt points out that a description of the place to act must uniquely specify the place, but it need not be the case that the one who hears the description know what the place is at the time the instruction is heard. It is sufficient if the hearer can be sure that there is a plan that will identify the referent when it becomes visible[5]. Appelt explains the joke where one passenger riding a bus asks another at which stop she should get off, and gets the reply "one stop before I do" as a case of description which does uniquely specify a location, but for which there is no effective plan to locate.

The Back Seat Driver uses a landmark when it can, and otherwise gives a distance. Landmarks are described in the next section.

A numeric distance is the cue of last resort. Very few natural direction givers used distances, and almost all of my test drivers complained about difficulty understanding distance. The voice directions tested by Streeter[82] included distance, even though they are not much use, because it could be easily and accurately calculated, and did not hurt. I found that people differed in whether they wanted to hear distances. Some never wanted to hear them, others wanted them all the time¹, and others wanted them only when the distance exceeded a certain threshold. To allow for these differences, there is a parameter in the user model for the minimum distance expressed as a number. If the distance is below this, a qualitative phrase is used; if above, a number is used. The cutoff can be zero, in which case numbers are always used, or set to an infinite value, in which case they never are.

A cue is expressed either as a full sentence ("Drive to the end of the street, then ...") or a preposed preposition phrase ("At the next set of lights, ..."). Experience has shown that a cue should not be expressed by a preposition after the verb ("Take a left at the lights.") because some drivers start to take the left as soon as they hear the word "left". This may be because synthetic speech does not provide enough intonational cues for the driver to reliably predict the length of the sentence, leading the driver to act on syntactic information alone, and thus taking the sentence to be complete as soon as the word "left" is heard.

The description of a road feature depends upon whether or not it is visible. If it is, it can be referred to with a definite article ("the rotary", "the overpass"). If not, an indefinite article is used. The program cannot tell whether an entity is actually visible, so it uses distance as an approximation. If the feature is closer than one tenth of a mile, it is considered to be visible.

¹A subject who played a lot of football had a good sense for distance in yards.

A special case of cues is when the driver is *at* the place to act. When stopped a few meters from the intersection, it is wrong to say "Turn at the next lights" even if it is literally true. In some cases, drivers thought that this meant not the current set of lights, but some set further on. The Back Seat Driver thinks of itself as being *at* an intersection if it is less than thirty yards away, except that if there is a stop light at the intersection and the car is not moving, then the intersection distance is fifty yards, since cars might be backed up at such an intersection. When at an intersection, the Back Seat Driver says "Take a left here". In an earlier version, it said "now" instead of "here", but this evoked violent protest from drivers waiting for a traffic light.² People rightly resent being told to do something they have good reason not to do.

4.4.3 Landmarks

Traffic lights are very good landmarks because they are designed to be easily seen and drivers have an independent reason to watch for them, namely a desire to avoid accidents. In an earlier version of the Back Seat Driver I used "major intersection" as a landmark. A "major intersection" is an intersection with a street whose quality is "good" where there is also a traffic light. I found that drivers preferred directions that spoke only of the lights, since that is a more precise term. When referring to a traffic light, if the car is at the intersection for the lights, the Back Seat Driver uses a proximal deictic ("this" or "these", as opposed to the distal "that" or "those") to show it means the lights that are here.

Buildings can also be landmarks. The Back Seat Driver has a very small database of buildings as part of its directory of services. If it finds a building on

²Given that the system knows where it is relative to the intersection, and can make good guesses about where the light is, perhaps it is possible to build a computer vision system to look for the traffic light.

the corner, it includes it as a potential landmark; e.g. "Look for Merit Gas on the left side."

Other landmarks are features of the road, such as underpasses, bridges, tunnels, bends in the road, and railroad crossings. It is very hard to miss an underpass, since the overhead road fills the whole visual field for a time. Large bridges are likewise unmistakable, but small bridges might not be noticed. A weakness in the current map is that there is no distinction between the size of bridges. It is not at all clear that a short bridge over a railroad should be called a bridge, as opposed to an overpass. Perhaps a new category is needed.

Still another potential landmark is the road coming to an end. This is a landmark that is impossible to miss. Some people call such endings "T" turns. The Back Seat Driver says "Drive all the way to the end, then ..." A problem with this landmark is that people can recognize it too soon. Apparently some drivers take "the end" to mean not "the farthest you can go along this road" but just "the next intersection".

4.4.4 Street names as landmarks

A street name can be a landmark, but not a good one, unless the driver already knows the street. (Street names are appropriate for a driver who knows the street network, but not the route. In such a case, street names are the best kind of landmark, because they are meaningful miles in advance.) There are several reasons why street names should not be used. First, the driver may not hear the name correctly. Second, the driver may hear the name, but not know how to spell the name after hearing it, so she may not recognize the name in its printed form. This is especially a problem when the driver is from out of town. If you were told to turn left "Lemon-stir" street, would you recognize it spelled as "Leominster"?

Finally, even if the driver knows the spelling, street signs are often missing, turned around, or invisible due to weather or darkness.

The problem of hearing beats further discussion. Synthetic speech is less intelligible than human speech. This is a problem for everything the Back Seat Driver says, but the problem is worst in name recognition, because with other parts of the utterance the driver can use knowledge about the English language to help in interpreting the utterance. A driver who hears "bear left" as "bir left" can probably reconstruct the intended meaning. A similarly distorted name often can not be repaired. Worse, drivers may not only fail to understand, they may misunderstand. One driver misheard the names of several streets, yet was so confident of the correct hearing that he drove past the turn, despite being told to "take a right".

Despite all the problems that come with using street name, many drivers ask for them. To accommodate them, a parameter in the user model controls the use of names. If set, names are supplied as part of the instruction. When names are included, they are attached at the end of the instruction ("Take the second left. It's Elm Street.") rather than directly ("Take the second left onto Elm Street."), which hopefully weakens their salience somewhat and makes them more of a confirmatory cue than an essential one.

It can sometimes be difficult to form the proper name for a segment. The naming of streets is a difficult matter if the street is not an ordinary street. A bridge, tunnel, or highway should be named as a singular. We say, "the Halyard Bridge" and "the Sumner Tunnel". Access ramps are harder to name, because they must be described, not named. If the ramp leads to a superhighway, it can be called by the name of the highway. "Drive onto the Massachusetts Turnpike" is just as good as "Drive onto the ramp for the Massachusetts Turnpike", even if it is not literally true. If the ramp leads towards an obvious feature (an underpass, a bridge) it can be elided, and we can just say "drive towards the bridge". Otherwise,

the only possible name is "an access ramp", which is probably not much help, but at least says that the road has no name.

4.5 Advice

The Back Seat Driver gives advice about how to prepare for actions. There are two forms of advice, lane advice and speed advice.

Lane advice tells the driver which lane to get into (or stay out of) when applicable. The system gives lane advice as part of the instruction when approaching an intersection where it matters. The instruction may also include advice about what lane to be in *after* the intersection, in preparation for the next act.

Speed advice warns drivers to slow down if they are traveling too fast to safely negotiate a turn. The limiting factor for angular acceleration is the driver, not the cornering ability of the car. The average driver will accept no more than .1 G radial acceleration[57]. Radial acceleration is v^2/r , where r is the turning radius of the turn. The Back Seat Driver knows the geometry of the road, so it can predict the maximum tolerable velocity for the turn. It does not tell the driver about this speed (it assumes the driver will choose a comfortable speed without being told), but it does estimate the distance required to decelerate, and it tells the driver to begin slowing down early enough to do this gently.

4.6 Discourse

In order to generate more fluent text, the Back Seat Driver keeps track of what has been mentioned. Some instructions are obvious after having been given. If the system tells the driver to go straight through a set of lights, there is no

reason to repeat the instruction when actually at the lights. This is in contrast with a turn, where the driver hears advance instructions to know what to do, and immediate instructions to know when to do it. This can be important, for if the driver hears "go straight through the lights" twice, she may try to go through two sets of lights. Indeed, not only is it redundant to say it twice, it is misleading, because some drivers take this to mean that there is a second set of lights to be driven past.

To implement this, each instruction is able to determine whether it is obvious after having been given once. When it is time to speak the instruction, if the instruction has already been given, and it is obvious once spoken, then it is not spoken again.

The Back-Seat Driver also retains a history of the route. This allows it to generate cue phrases for the instructions. If the route calls for doing "the same thing" twice in a row (e.g. two successive rights), the system uses the word "another" to indicate this doubling. This is important for polite behavior. If I (as a passenger) give you (as a driver) instructions by just saying, e.g. "Take a right. Take a right. Take a left. Take a right.", pronouncing each the same, you will judge me to be rude. My speech is not acknowledging your actions or your history³. There are two ways I could acknowledge your work: I can use cue words, and say instead: "Take a right. Take another right. Now take a left.", or I can use intonation to indicate that my instructions are all part of a series. In this case, I would use the so-called "list" contour[49]. The Dectalk speech synthesizer does not support flexible control of intonation, so cue words are the only possibility.

4.7 Mistakes

If the driver leaves the route the system immediately informs the driver and

³I thank Chris Schmajuk for this explanation.

begins to plan a new route. Route planning after a mistake is no different from any other time, except that the vehicle is more likely to be moving. Telling the driver what she did wrong prepares her for hearing new instructions, and perhaps helps her learn to better interpret the style of language that the Back Seat Driver uses.

To describe an error, the system looks forward from the last saved position for the first action, even if obvious. (While the driver is on a route, the system records her position, so that if a mistake occurs the system can find the place where the mistake occurred.) This is action that the driver failed to perform. It utters a description of this action, saying e.g. "Oops, I meant for you to take a right." This error message is a change from the original message, which was of the form "You made a mistake. You should have taken a right." The new message is neutral, because test drivers complained about being blamed. A driver might leave the route deliberately, or the error could be system's, not the drivers.

An unsolved problem is how to reliably detect when the driver has left the route. At present, if the driver turns into, say, a gas station, the system will believe, falsely, that the driver has turned onto some street, because the street map includes only streets, and not other paved areas such as parking lots and filling stations. From this false belief, the system will conclude that the driver has made a mistake. If the street map included more information, the system could avoid this error.

Sometimes the driver will choose to not follow for good reasons that the Back Seat Driver is unaware of, perhaps because the road is blocked or because of a traffic jam⁴. In the first case, the driver can push a button - the "I Can't Do It" button - which informs the system that the road is (temporarily) blocked. The system automatically finds a new route. In the second case, the driver's only

⁴In the future, we can expect that the Back Seat Driver will be monitor machine-readable broadcasts to obtain such information, and plan a new route automatically.

recourse is to cancel the current trip (by pushing another button), then drive in the desired direction, and then request a route to the original destination (which she may easily do by touching a single button). It is essential, though, that the driver either notify the Back Seat Driver of the impossibility of the requested action, or cancel the trip, because otherwise the system will treat the deviation from the route as a mistake, and continue to attempt to find a new route, which may very well lead back through the street the driver is trying to avoid.

4.8 Reassuring

While the driver is following a route, the system adopts a persistent goal of keeping the user reassured about her progress and the system's reliability. If the Back Seat Driver were a human, this might be unnecessary, since the driver could see for herself whether the navigator was awake and attending to the road and driver. But the driver can not see the Back Seat Driver and so needs some periodic evidence that the system is still there.

One piece of evidence is the safety warnings the system gives (page 90). But if all is going well, there will not be any. The system gives two other kinds of evidence that things are going well. First, when the user completes an action, the system acknowledges the driver's correct action, saying something like "nice work" or "good". This feature is very popular with most test drivers. When I turn it off, most people complain. A few find the choice of words patronizing, but still want the function. Those who want no confirmation messages can turn them off by changing the user model.

The second form of evidence is to make insignificant remarks about the roads nearby, the weather, and so on. If the driver assumes that the navigator is being cooperative, as set out in Grice's maxims of cooperative conversation [28], then the

driver can infer that everything is going well, for otherwise the navigator would not speak of trivial matters.

Chapter 5

User and System Goals

The Back Seat Driver is a truly *interactive* system, by which I mean that both parties are active all the time. Most "interactive" computer systems are better called *reactive* - the machine and the human take turns. A typical reactive system is the command interpreter for a time shared operating system. While the human is typing, the computer does nothing except echo typing. It begins to work only when the human hits the Return key. Then it's the human's turn to wait. While the machine is working, there is little the human can do, except to interrupt a calculation gone awry. Giving driving instructions is an inherently interactive task because the concept of "turn taking" does not apply. The driver is always in some kind of action. The driver may be waiting (e.g. for a traffic light to change) but is never (or should never be) waiting for the system to say what to do. The task of driving does not divide into discrete steps, with an arbitrary pause between each step. The driver may choose to stop driving for a time, but if the driver is driving at all, she's driving all the time. So the machine must be ready to speak when needed.

5.1 Resource allocation

The Back Seat Driver pursues many goals at once. Its main goal is to get the driver to some location. This goal requires many actions over an extended time. Other goals include delivering messages, weather reports, and notifications¹. These goals can not all be met at once, so it must choose which to pursue at any given moment. To make this choice properly requires a model of action which accounts for the time required to perform an action. The only kind of action the Back Seat Driver takes is talking. Talking takes time, and only has value if completely performed. Talking is different from, say, drinking a cup of coffee. A partially drunk cup still has some effect. An incomplete sentence has an effect, too, but not a useful one. The program should only take an action when it can expect to complete it.

Since the Back Seat Driver can only say one word at a time, it must choose its words carefully. At any moment when the Back Seat Driver is not already speaking, it can either begin saying something, or it can wait. This decision must be made again and again, at each passing moment. A wrong decision can not be undone. Words once spoken can not be made unsaid, though the program can stop talking at any time; more seriously, there is nothing to be done if the program decides at some moment that it should have begun speaking two seconds previously². Since the driver's actions are unpredictable, the program must improvise. Plans can not be made in advance.

The problem, then, is to allocate resources to goals so that the most important goals get all the resources they need, while lesser goals get any useful amount left over. Although the only resource used is speaking time, the method used is designed to be applicable to goals that use many different kinds of resources.

¹ and someday, reading mail, advertisements, tourist information, and traffic advisories

² Save perhaps to speak faster, a possibility I do not consider further.

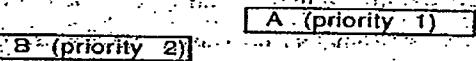
Each goal provides three pieces of information about itself:

- Is it ready to run - (e.g., is there something that could be said right now which would help achieve the goal?)
- If so, the resources it will require, and the maximum time it will require each resource.
- If not, the minimum time until the goal will be ready to run.

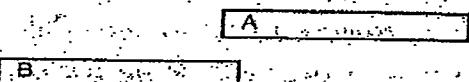
Each goal decides for itself whether it is ready to run. The criteria depend upon the goal. For instance, the goal to deliver electronic mail will only be able to run if there is some mail to speak. A goal can be certain of whether it is ready to run, but the latter two measures may be estimates if they depend on the driver's future actions. For instance, while following a route, the time when the next instruction will be given depends on the speed of the car. If the driver speeds up, the time will come sooner, and if the driver makes a mistake, corrective action is required at once. Given these estimates, and a preassigned priority for each goal, the program will at each moment authorize at most one goal to run. It does this by examining each goal in order of decreasing priority. The goal of highest priority that is ready to run will be allowed to do so, unless there is some other goal of greater importance which uses a resource used by the lesser goal and, though not currently ready, expects to begin in less time than the lesser goal will use in speaking. (See figure 5-1.) Thus the Back Seat Driver does not start a narrative it can't finish. Nevertheless, interruptions do sometimes happen when driver does something unexpected or adds a new goal.

5.2 Architecture

This section discusses the architecture of the goal mechanism.



(a) Goal B has time to finish before Goal A starts.



(b) Goal B is not allowed to start

Figure 5-1: Resource allocation

The unit of planning in Back Seat Driver is the goal. Every goal has a name and an importance, which is a number from 1 to 9. Goals may also have any number of slots, each of which has a name and a value. Slots hold information about the goal. The unit of execution is the goal function. A goal function is a module capable of taking action on behalf of a goal. The definition of a goal function includes the name of the goal it can work for. In theory, there could be several goal functions relevant to a goal, but at present there is always exactly one. A goal is mostly a holder for slot information. The actual work is done by the goal function. In the discussion that follows, the term "goal" is used for both goals themselves, and goal functions, except where the distinction is important. The unit of resource allocation is the resource. There is only one resource, the speech resource.

The system maintains a list of goals. When a goal is added to the list, the system finds the appropriate goal function and attaches it to the goal. The goal function can then access the slots of the goal to determine what it needs to do.

The system communicates with goals by message passing. The goal function protocol includes the resource allocation messages described on page 84 and a

message which allows a goal to run. When a goal runs, it is expected to take one action, and then return. The system has no way of interrupting a goal function while it is running, so actions should be quick.

Most goals achieve their aims by speaking. A goal speaks by requesting the speech resource to send characters to the speech synthesizer. This can be done fairly quickly, but the synthesizer requires substantial time to finish speaking. This is why I say that actions take time. During this time, the goal can make no further progress. The goal must be able to tell when speech is complete, for only then can it claim to have achieved the desired effect. (I assume that if the speech was spoken, it was also heard and understood by the driver. If not, the driver will take some action such as requesting a repeat.) The speech resource retains the identity of the goal which last spoke. Periodically it polls the synthesizer to determine whether the speech is complete. When the speech has finished the speech resource becomes available again. Usually the goal that produced that speech will run a second time, notice that the speech is complete, and finish its operation.

Goals may interrupt lower priority goals by requesting the speech resource to interrupt the lower priority goal. Interruption stops the speech synthesizer immediately. The interrupted goal is informed of the interruption, and can react as it chooses. There is no way for the goal to know whether any of its words were actually spoken, so it has to start all over. Most goals attempt to run again as soon as possible. The "help" goal simply quits when interrupted, on the assumption that the user has learned whatever she wanted to know, so no further help is required. The assumption is that the interruption occurred because the user started some higher priority goal after learning how to do so through the help command.

Not all actions of a goal are of equal importance. A goal can lower its own priority (but not raise it) if it wishes to say something of lower importance.

Goals may have subgoals. Subgoals can be executed for side-effect or to return a value. (The route finding goal returns the route found as a value.) A goal with a

subgoal is not eligible to run until the subgoal completes. A goal that is a subgoal has a slot pointing to its parent goal. When a goal finishes, if it is successful, it adds a slot to itself to hold its returned value. This slot is empty if no value is returned, or if the goal failed. It then informs the parent goal of the outcome (success or failure). If a subgoal succeeds, the parent goal is again eligible to run. It can get the value returned by the subgoal by looking at that goal's result slot. If the subgoal fails, the parent goal may try again or may give up.

5.3 Kinds of Goals

Goals can be temporary or persistent. Temporary goals can be satisfied, but persistent goals never can be. Temporary goals are about attaining a condition which does not currently hold, and persistent goals are about maintaining a condition which is true now, should stay true, but might become false without at least occasional action by the system.

The user creates a goal by striking a button on the cellular phone keypad. (In a more advanced system, this might be done through speech recognition.) Examples of user initiated goals are:

- to get to some location
- to find the closest provider of a service
- to hear a weather report
- to hear the last statement repeated

In order to get to a location, the system must first know what the location is and how to get there. For each of these it creates a subgoal. The subgoal to get the address of the desired location asks the user to enter it with the cellular phone.

keypad, in the same manner used in Direction Assistance, (see page 102). The subgoal for finding a route invokes the route finder, as discussed in Chapters B and 6. Once the system has a route, it creates two subgoals, one to have the user follow the route, and one to keep the user confident of her progress along the route. These goals are described in chapter 4. There is just one goal for following the route, rather than a separate subgoal for each intersection along the route. There is no large theoretical reason for this, but it does make it easier to give information about the route as a whole (e.g. "How much longer is it?")

5.3.1 Finding Services

Drivers sometimes know what they want, but not where to go to get it. What they need is a geographically coded directory of services, like the Yellow Pages. Drivers may want to browse the directory and select a service themselves. If they are interested in a commodity where the vendor is unimportant, it may be more convenient to let the system select the destination. When running low on gas late at night in an unfamiliar area of town, you want to find the closest gas station that is open. You do not care which station you go to, you just want the closest. The Back-Seat Driver can find the closest provider of a service. The directory of services, an electronic "Yellow Pages" has entries for services indexed under topics. The directory holds only three kinds of services: gas stations, automated teller machines, and ice cream stores. The entry for each service includes the name, the address, the hours of business, and a phone number. Upon the request³, the

³The present interface will not extend gracefully to a much larger number of services, because it asks the user to select the kind of service by choosing from a list in the same way that that the user might select a city when choosing among alternative destinations. The Boston Yellow Pages has more than 3600 categories, so choosing in this way would take about an hour on the average. On the other hand, it is unlikely that a "spelling" interface would be successful, since there are many alternative wordings for each category. Selecting from a list this large is better done with a textual interface which allows rapid scanning.

system finds the geographically closest⁴ provider known to be open, and offers to find a route to it.

5.3.2 Other user goals

The system obtains weather information from an online weather station. The station returns a structured list of the current weather conditions, and the Back Seat Driver assembles a set of sentences to speak the information.

The system treats "repeat the last statement" as a goal, rather than as a special purpose function, except that the importance of this goal is set to the value of the the last goal to speak (the goal whose utterance is being repeated). This guarantees that if some more important goal desires to speak, it will be able to. A repetition of an utterance is no more important than it was originally.

Repetitions should not necessarily be literal. If the situation has changed, a repetition of what was once true may now be a lie. For example, the car may have moved, and the turn that was once the "second right" might now be the next right. The goal function for repetition determines the last goal to speak by asking the speech resource. This goal is then asked to generate an utterance equivalent to its previous utterance, providing for the opportunity for the goal to choose different words (or a wholly different message) if need be.

5.3.3 System goals

All system initiated goals are persistent. The system goals are:

- to inform the user of new electronic mail

⁴It would be better to find the one that is most easily reached, but it is too expensive at present to compute.

EXHIBIT

2

(Part 6)

- to deliver messages from the base station.
- to warn the driver of dangers ahead.

These goals can never be satisfied: mail or messages can arrive at any time, so the user's safety should always be preserved.

The first two goals are easily implemented. The Back Seat Driver can check for new mail by consulting the computer that stores electronic mail. This computer is accessible via a local network. At present, the Back Seat Driver can only say that new mail exists. It would be possible, and desirable, to read the mail aloud, and to integrate voice-mail and text-mail, as in the Phone Slave[75, 77]. The Back Seat Driver can also deliver messages entered at the base station. Communication from the laboratory to the car has mostly been useful for debugging (and the occasional joke).

The Back Seat Driver can warn the driver about those dangers which can be inferred from knowledge of the road network. These dangers include:

- driving above the speed limit
- driving the wrong way on a one-way street
- driving too fast for an upcoming curve
- driving on a one-way street that becomes two-way ahead
- merging traffic
- "blind" driveways ahead
- speed traps
- poorly repaired roads
- dangerous intersections

These dangers are not of equal severity. Some are potentially fatal, others mere annoyances. The first five have been implemented. The others can be added if the required features were added to the database. A difference between the Back Seat Driver and other systems which have provided safety warnings is the Back Seat Driver attempts to determine hazards by reasoning about road conditions rather than requiring them to be built in. Warnings are remembered, and not repeated.

Speed limits and one-way streets are included in the map data base. If the driver is driving more than five miles per hour above the speed limit, the system issues a warning. This feature was actually requested by some users. Warnings about driving too fast were not nearly so popular. The maximum safe speed is computed by finding the safe turning speed (see page 77) for each possible segment at the next intersection. To avoid false alarms, the maximum value is taken. This speed is too great if the driver intends to turn, but it is better to give false negatives (i.e. be silent) than false positives.

The Back Seat Driver warns drivers about potential dangers in traffic flow. There are intersections in the Boston area where traffic flows from a one-way street straight onto a two-way street. The danger here is that a driver could take the oncoming left lane to be a second lane for forward travel. A second "flow" problem is where the number of lanes is reduced ahead. The driver may need to change lanes, or watch for other cars merging in at the last minute. Both these dangers are marked with warning signs on the street, so it appears reasonable for the Back Seat Driver to also give warning.

Giving warnings is itself something of a danger. One person who tested Back Seat Driver remarked that the system did such a good job of telling her what to do that she assumed it would watch over *all* details of the driving, and she did not feel as responsible herself. It is not at all clear what to do about this possibility.

Chapter 6

Comparing Routes

Drivers want the Back Seat Driver to find the best route to their goal, or at least a good one. This means that the Back Seat Driver needs a way of comparing two routes, and selecting the better of them. This chapter describes how the Back Seat Driver makes such comparisons. In addition, since different drivers have a different sense of what makes a route "good", the Back Seat Driver has three different comparison functions. The driver can define a good route as one that is either short, or fast, or simple. (Planned, but not implemented, are the abilities to find "scenic" routes, novel routes, and routes that pass near providers of services, for example gas stations.) Readers who are not already familiar with the A* search algorithm should consult chapter B before reading this one.

6.1 Comparing routes requires a metric

To find the best route requires a function which can compare two routes and select the better of them, according to some fixed criteria provided by the user. For

technical reasons described in chapter B, the comparison function must not only select the better of two routes, it must also provide a numeric rating of "goodness" for each route. Such a function is called a metric.

The Back Seat Driver has three different metrics. The distance metric finds the shortest route, the speed finds the fastest route, and the ease finds the easiest route. Figure 6-1 shows the route each metric produces for a trip from 12 Albany Street to 4 Glenville Terrace. Compare these routes with the route provided by a human in figure 2-2.

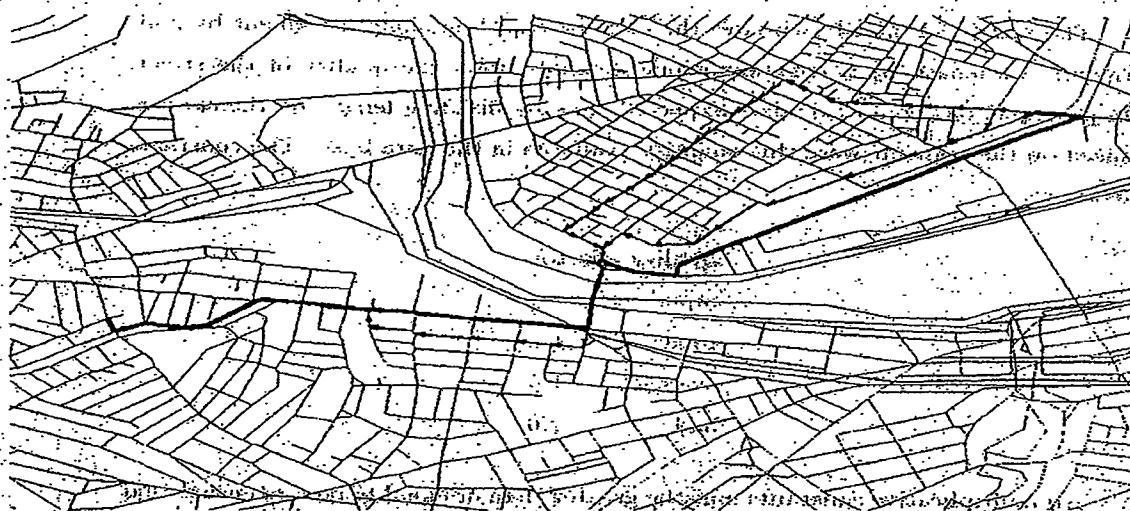


Figure 6-1: Four routes from 12 Albany Street to 4 Glenville Terrace

The shortest path in figure 6-1 is 2.73 miles long. Neither the "fastest" nor "easiest" paths found are optimal by their respective metrics, because the A* bias parameter (see page 147) is greater than 1. This makes the route finder less willing to go out of the way in searching for a route. The actual "fastest" path is shown as a thick line. This is also the "easiest" path. When finding the shortest path,

the A* bias parameter (D , see page 147) is 1. This makes search slower, but if the driver is really interested in the shortest path, the extra time to search is justified.

The distance metric is just the sum of the lengths of the component segments. The other two metrics are more complicated than the distance metric because they must consider intersections as well as segments. In general, there is a cost to travel along a segment and a cost to get from one segment to another. All costs are expressed as an "equivalent distance" which is the extra distance one would travel to avoid the cost.

The metric for speed estimates the cost for traveling along a segment by multiplying its length by a constant which depends upon the quality of the street. In principle, one could calculate expected time by dividing length by the average speed on the segment were this quantity available in the data base. The constants used are:

	Quality factor
super	1
good	1.2
average	1.5
bad	2.0

All multiplicative constants must be greater than or equal to one, to ensure that the cost of a route is never less than the straight line distance between two points. This condition is essential to the correct operation of the A* search algorithm, since the estimation function (g^*) must always return an under-estimate.

The time to cross an intersection is modeled by a mileage penalty which depends upon the nature of the intersection.

Factor cost reason

turn 1/8 mile Must slow down to turn

left turn 1/8 mile May have to wait for turn across traffic flow

traffic light 1/8 mile Might be red

If the segment is one-way, the penalties are cut in half, since there will be no opposing traffic flow. The turning penalties are computed based only on the angle between two segments, not on segment type or quality. The angle used is the smallest possible angle for the two segments, given the uncertainty of positions inherent in the DIMB file. This nullifies the DIMB "wobble" described on page 46.

The metric for simplicity seeks to minimize the driver's effort in following the route. Again, driver's effort is the sum of the effort to travel along a segment and the effort to get from one segment to another. Travel along a segment depends upon its quality. A limited access highway is the easiest to travel along, since all traffic is going the same way and there are no intersections. Low quality segments are the worst to travel along - the driver must be wary of potholes, and their narrowness requires careful attention. Turns of every sort are penalized equally, since they all require decisions. The intention of this metric is to find routes which require the least amount of speaking by the Back Seat Driver, leaving the driver free to concentrate on other matters. A more advanced Back Seat Driver should perhaps select this mode automatically if the driver is listening to music or having a conversation.

6.2 Estimating the time required to find a route

The Back Seat Driver must be able to find routes while the vehicle is moving, yet the route finding algorithm requires a fixed origin. When the car is moving, the Back Seat Driver first estimates the distance the car will travel during the

route finding process by multiplying the current velocity by the estimated time to find the route. Then it finds the position the driver will reach after traveling this distance, assuming that the driver will not make any turns without being told to do so. It then finds a route from this extrapolated position to the goal. Finally, it finds a route from the cars actual position to the estimated starting position. This second route is so short that the car is unlikely to move far during the time it is computed.

The route finder estimates the time to find the route between two points by multiplying the distance between them by a constant. This constant was initially determined by running the route finder for 20 randomly selected pairs of origins and destinations. As the Back Seat Driver runs, it accumulates additional values for the constant.

Part III

Conclusions

Chapter 7

Related Work

This chapter first surveys related work on computer programs which provide navigation assistance to drivers, then develops a taxonomy of such systems. A related survey can be found in [56].

7.1 Early Work

Early application of computers to navigation was intended to reduce traffic congestion by providing route information to drivers. The designers of the Electronic Route Guidance System (ERGS) intended to make traffic flow more efficient by balancing load. They believed that reducing driver uncertainty at decision points would make traffic flow faster and more safely. In the ERGS design, a driver beginning a route finds the intersection closest to the destination, then enters a five letter code word for the intersection. When the vehicle passes over an induction loop sensor in the road it transmits the destination to a central computer. The computer determines the best route, and relays instructions to the car. This

interchange of information occurs at every instrumented intersection. Driving directions combine simple text from a nine-word vocabulary and directional arrows, and are displayed by a "heads-up" display. The ERGS system was designed, but never implemented[70]. A similar system was designed and tested in Germany in the late seventies[14].

7.2 Elliot and Lesk

The pioneering work on computer navigation assistance is by Elliot and Lesk[21, 22]. They showed that general purpose graph search algorithms¹ are not suited to the problem of finding useful routes for people traveling in the real world. There are two reasons why this is so. First, general purpose graph search algorithms are complicated because they must work with any kind of abstract graph; but street maps are among the simpler forms of abstract graphs. The extra complexity in an algorithm for arbitrary general graphs makes it slower than one which is specialized for searching simpler graphs.

A second problem with general search algorithms is that the shortest route may not be the *best* route. It might be a maze of shortcuts. Elliot and Lesk say that people who saw such routes "recoiled in horror", and so they modified their algorithm to prefer a route which was slightly longer but had fewer turns. The usual algorithms for graph search consider only the cost of traveling along an arc.

¹Abstract graphs are the mathematical basis behind all computer route finding. An abstract graph is a set of nodes (points) and arcs (lines joining two nodes). To represent a street map as a graph, think of intersections as nodes, and streets as arcs. If two intersections are directly connected there will be an arc between the corresponding nodes. A route between two nodes is a chain of arcs, each leading from one node to the next, such that the origin node is the first node, and the destination is the last node. Each arc may have an associated cost (greater than zero) which represents the effort required to travel the associated street. Usually this would be the length of the street, but it might also represent the toll collected on a turnpike. If the cost is the distance, then the shortest route is the route whose sum of whose arcs is least. Finding a route, or the shortest route, through a street map is a special case - and perhaps the best example - of the problem of searching an abstract graph.

(a street), and take the cost of going from arc to arc to be zero—that is, having arrived at a node, all arcs are equally accessible. But this is not true of street driving. It takes more time to make a left-hand turn across oncoming traffic than it does to go straight, and it takes extra effort to locate the place to turn. Elliot and Lesk added a system of weights, which are extra costs associated with turns. To their algorithm, the cost of a route was the sum of the distance along the streets and the difficulty of making the turns. They set the weight of a left turn to be 1/4 mile and a right turn to be 1/8 mile. This caused the algorithm to prefer slightly longer routes with fewer turns to short, twisty routes. Every known route finder trades distance for simplicity.

Elliot and Lesk also were the first to implement a program to generate natural language driving instructions for the route. This is not a straightforward translation. To understand why, you must understand that the map represents each street as a set of segments, where a segment is a piece of a street chosen to be short enough that it is a straight line and no intersection with any other segment occurs except at a segment endpoint. The route as represented by the route-finding algorithm is a sequence of street segments, not of streets. A street segment does not match any common-sense notion of a road. Route descriptions must be expressed in terms of motion along streets (across many segments) and turns, not as a list of segments.

In their instructions, a route consists of a beginning, a sequence of turns and crossings (of rivers or railroads), and an ending. For each of these, there is a template to generate a sentence. A template is a sequence of words and slots representing fixed and variable components of a sentence for a given type of act. The words are copied directly to the output, and the slots are filled in according to the particulars of an act. An example template is:

```
Go <distance> [<intersections>] turn <direction> on <street>
```

Here "Go", "turn", and "on" are the fixed words, and everything else is a slot. The slot <intersections> is optional. This template might produce:

Go 0.3 miles (2 intersections) turn left on TROY HILLS RD.

A third contribution of Elliot and Lesk was to integrate the digital map with other location oriented databases, including a Yellow Pages and a personal address book. This allowed the program to find routes to addresses given a person's name, to find the closest store of a specified category, and to mention stores along the route as possible landmarks.

7.3 Direction Assistance

Direction Assistance[19] provides spoken directions between locations in the Boston area. It uses a Dectalk speech synthesizer. This synthesizer includes a telephone interface, so it can answer a phone call and decode touch tone button presses. To use Direction Assistance, you call it from a touch tone phone. It answers the call, and prompts you to enter your origin and destination locations as street addresses. It finds a route, then describes the route to you. Direction Assistance was directly inspired by Elliot and Lesk, and extends their work in three ways:

The most significant difference is that Direction Assistance speaks its directions, where Elliot and Lesk drew maps and provided written text. Using speech makes the program much more accessible, since only a touch tone phone is required, rather than a computer terminal. The disadvantage is that users must remember the instructions or write them down.

7.3.1 Entering addresses

The method of address entry is of some interest. You enter an address by first entering the digits, then a number sign, then spelling the street name using the letters on the telephone keypad. The letters "Q" and "Z" are on the "6" and "9" keys, respectively, and the space character is on "1", which is otherwise unused. These keys are sufficient to spell any street name in Boston. (The spelling rules would require expansion to enter street names with other characters in them, for example "4th Street".)

Spelling a street name requires only one button push for each letter, even though there are three letters on each key. This is because of the redundancy in street names, which are pronounceable words, not arbitrary strings. There are 37 pairs of street names with the same "spelling" in the reduced "alphabet". An example is "Flint" and "Eliot", both encoded as "35468". This is only one percent of the (2628) names of streets in Boston, so collisions are rare. This technique appears workable even for larger sets of names. When the entire word list of the Brown corpus are encoded, there are still only 1095 collisions in the 19837 words (5.5%).

If a name collision occurs, the interface reads the list of possibilities and asks the driver which one was meant. This is very rare. A more common problem is that street names are duplicated. For example, there are 14 possible meanings for "10 Washington".

- 10 Washington Avenue, East Boston
- 10 Washington Avenue, Union Square, Somerville
- 10 Washington Avenue, Chelsea
- 10 Washington Court, Central Square, Cambridge
- 10 Washington Place, South Boston
- 10 Washington Square, Charlestown
- 10 Washington Street, Charlestown
- 10 Washington Street, Central Square

- 10 Washington Street, North End, Boston
- 10 Washington Street, Union Square, Somerville
- 10 Washington Street, Brighton
- 10 Washington Street, Brookline
- 10 Washington Terrace, Union Square, Somerville
- 10 Washington Terrace, Charlestown

(This is the worst case example. There are only half as many possibilities for "100 Washington".) When this happens, the system asks the user a series of questions to reduce the list to a single choice. The system tries to ask the fewest questions possible. It asks the user to choose from a list of street types; if that is sufficient to resolve the question, and otherwise from a list of the containing cities (or neighborhoods, if there are two instances within a single city). To select from a list, the system reads the contents, asking the user to push a button when the desired choice is read. The interface is described more fully in [18].

A problem not addressed by Direction Assistance is that some "addresses" do not refer to streets at all, but rather are the names of buildings or developments, e.g. "11 Cambridge Center" or "One Kendall Square". Direction Assistance can only understand addresses expressed in terms of named streets.

7.3.2 Generating text

A second significant difference between Direction Assistance and the work of Elliot and Lesk is that Direction Assistance generates better quality descriptions of the route. The improvement arises because the text generation process first analyzes the route into a sequence of "acts", and then generates descriptions from these acts, instead of working directly from the route. An act represents something that the driver does. There are eleven different acts, each representing a different way of moving. The type of act depends upon topology (how many streets are present at an intersection, and which way traffic can flow), geometry (what angles

the streets make) and what kind of streets are involved. Thus we say "bear right at the fork" rather than "turn right"; but we don't say that in taking an exit from a highway we are "bearing right". An act may involve more than one segment, as for instance a "U Turn" on Mémorial Drive (shown in figure 7-1) takes one from Memorial Drive, to Danforth, and back onto Memorial Drive, yet should not be described as two successive turns. For each act there is a specialized text generator

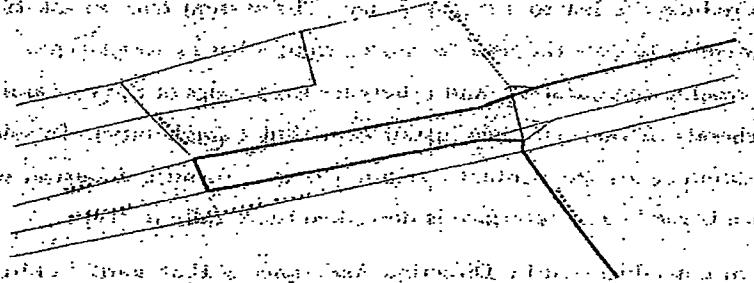


Figure 7-1: A U turn

to describe it. There is also a function to find an appropriate cue or landmark (e.g. a street crossed or an underpass) just before the location of the act.

7.3.3. Route Finding differences

The Direction Assistance route finder uses a different algorithm than Elliot and Leski, and has a different set of weights. The algorithm is the A* search[35]. The weighting scheme ranks roads by a four-valued "goodness" feature and penalizes routes that use less good roads by multiplying the mileage by a constant factor. It also reduces or waives the penalty for turning under a set of circumstance having to do with predicted ease of following; for example, a turn onto a one way street incurs a lesser penalty, since it is unlikely that the driver would turn the wrong

way. It reduces the penalty for "T" turns since the driver can not possibly miss the place to turn. In practice, these weight changes have very little effect.

7.4 Text-Based Directions

7.4.1 Counter-Top Directions

The Hertz car rental company offers "Computerized Driving Directions" at some of its rental counters. The directions include approximate mileage and estimated travel time, but are highly schematic, even cryptic. An example appears in figure 7-2. Despite appearances, these instructions are not computer generated.

```
APPROXIMATELY 16.8 MILES 0 :35 TIME
2.0 MI NORTH TO I-78 WEST enter LEFT
14.90 MI WEST TO NEW PRO/BERKELEY EXT bear RIGHT
DIAMOND HILL RD continue
0.4 MI TO MOUNTAIN AVE turn RIGHT
0.4 MI TO AT&T/BELL LABS on your right
```

Figure 7-2: example of driving instructions provided by Hertz

The Hertz system is more akin to a database retrieval system than a route finder. A California firm, Navigation Technology, sells a product called "Drive Guide" which is reported to be able to print driving directions between any two points [71, 4] in the San Francisco area.

7.4.2 Ma

Peeder Ma describes a system which gives textual directions in [51]. His work is similar to both Elliot and Lesk's and to Direction Assistance, but was apparently created independently of both. Ma uses A* search with a penalty factor to

minimize the number of turns. Unlike Elliott and Lesk, he uses the same penalty for both left and right turns. His street map representation does not include one-way streets or restrictions on turning ("no left turn") so it does not always find usable routes. His route descriptions use a taxonomy about as elaborate as that of Direction Assistance, but the text generated is more stylized.

7.5 Automotive Navigation Systems

Several groups have built position or navigation systems for use in automobiles. For the most part, these systems have not been well described in the literature, probably from a desire to preserve commercial secrecy.

The most well known in the United States is the ETAK Navigator, which displays the car's position on a map display on the dashboard [36, 89]. The map rotates as the car turns so that the forward direction is always straight up on the map. The ETAK Navigator can show the map at four different scales. At greater scales, the display shows only large streets in a straightened form. This is important to keep the map legible at large scales. The system provides a limited amount of navigation assistance. The driver may enter a destination (as a street address or intersection), and the system will display the direction and distance to the point. It remains the driver's task to select an appropriate route to the destination.

The Routerechner provided directions in and between German cities [30]. The route finder could receive real time traffic information by digital radio while on route. This system's map included only the Autobahn, and not the cities (this was before CD-ROMs were widely available), yet it was also provided a limited navigation service within cities. The user entered the destination as a pair of coordinates, and the system displayed the direction and distance to the destination. As with

ETAK, it was the user's responsibility to select an appropriate road. Haeussermann reports that users were always successful at finding their destinations and were pleased with the system.

The Honda Electro Gyro Cator[83] provided displayed position of the car by plotting a point on a screen. The driver could determine position by placing a transparent map over the screen. This system did not provide route directions.

The Nissan-Hitachi car navigation[38] and information system displays position on a map, finds the shortest route to a destination taking into account real time traffic information, and gives directions by arrows on the face of a display. The system also includes a "secretary mode" which displays the driver's appointments. The system uses a CDROM for map data, and combines satellite positioning with dead reckoning.

The EVA[65, 58](Automatic Navigation System) developed in Germany by Blaupunkt, the University of Karlsruhe, and the federal government, accepts destinations as street addresses, finds minimum time routes, and gives directions by a combination of simple (arrow) graphics and voice. The system can recover from a driver error in following the route and find a new route within 50 meters of travel.

The Phillips corporation, in the Netherlands, is developing a prototype car information and navigation system called CARIN[85, 8]. The driver enters a destination using either a keyboard or a touch sensitive screen. The system displays routes on a map and gives spoken driving instructions. The map is stored on board in CDROM, and a radio link provides for updates on traffic conditions. The system is potentially interesting, but very little has been published about it.

7.6 Classifying navigation systems

Navigation systems can be classified along three dimensions. There are three

kinds of navigation service:

- positional systems tell you where you are.
- orienting systems show the direction of your destination.
- instructional systems tell you what to do.

A navigation system can provide one, two, or all of these services. Navigation systems can be further distinguished by how they provide the information:

- verbal systems speak.
- text systems provide text.
- graphic systems provide pictures.

Finally, systems can be classified as either real time or static. The categories of this classification are not independent. There can be no static positioning system, since one can not predict the future position of the car.

The systems of Elliot and Lesk, Ma, and Hertz provide static, text instructions. Direction Assistance gives static verbal instructions. There are several problems with "static directional" navigation systems. First, they do nothing to help the driver follow the route. The driver must determine for herself when to apply each instruction. Instructions like "drive half a mile, then turn left onto Maple Street" are no use if the driver is unable to measure mileage or can not determine the name of the street. The urban street network contains many short connecting roads (access ramps) which are nameless. Finally, even a named street might be missing its sign. In addition, the driver must keep track of which instruction is next. A second problem is that since the instructions must be specified in advance, there is little to be done if the driver does not follow the instructions, which might happen from error, or because the instructions are wrong, or simply ill-advised (as

when confronting a traffic jam). The simulated navigation system constructed by Streeter included instructions of the form "If you see this you've gone too far" but none of the actually implemented system did this:

Most systems which provide positions also provide orientation. (The Honda Gyro-cater provided position only, and the Routerechner provided orientation only). Positional or orientational systems can be useful for navigation provided one can read a map and find one's own route. It is not clear whether the CARIN system retains a map as a "vestigial" display, or because its makers do not appreciate the superiority of speech, or because they see a need for positional information other than route finding.

Chapter 8

Future Developments

The Back Seat Driver works. It works well enough that one can imagine something very much like it being sold within a few years. There remain some areas for further research. Some of these involve only the Back Seat Driver, while others have to do with how a system like the Back Seat Driver would fit into the larger context of urban planning and public policy. The Back Seat Driver of the future will not simply be something installed in a car, rather it will be part of a network of information and services that includes your home and office, other drivers, and perhaps the local police agency.

This chapter discusses areas for further research. The topics are arranged roughly in order of increasing scale, ranging from the connection between the Back Seat Driver and other systems in the car to the connections with Federal law.

8.1 Integration with the car

The Back Seat Driver should be running on a computer embedded in the car, so that it can get more and better information about the state of the car and driver.

For instance, when the next instruction is a turn, the Back Seat Driver should notice whether and when the driver turns on the turn signals. If the driver applies them too soon, it is possible (but not certain) that the driver has underestimated the distance to the turn; if applied at the "right time" then the system can take that action as confirmation that the instruction has been understood; if never applied, then the driver has either misunderstood, or is driving hazardously!

The Back Seat Driver should also be integrated into the car's audio system, rather than having separate systems for voice and music. Furthermore, it should pay attention to what the driver is listening to. If the driver is listening to the radio, or playing a CD (or using a cellular telephone) the program should try to speak less often, on the grounds that the driver has implicitly indicated a preference for what to listen to. The program should suppress reminders and historical notes altogether. When it must speak, it should borrow the audio channel rather than trying to speak over it. The Back Seat Driver should also be aware of the driver's use of other controls in the car. It should defer speech while the driver is adjusting, e.g. the heat or the mirrors, and suppress speech altogether if the car makes sudden extreme changes in velocity. A driver trying to cope with an emergency situation does not need another distraction.

Back Seat Driver should also learn about the performance of the car. A suitably instrumented car could also measure the coefficient of friction by comparing the applied braking force and the resulting deceleration. This would allow it to adjust the time factors used in deciding when to speak.

8.2 Knowing the driver

While it may be so that "all men [meaning persons] are created equal", it is not
for the signals may be broken.

the case that all persons should be treated alike. Car seats, mirrors, and suspensions are adjustable now, and the Back Seat Driver should also be adjustable, or better yet, it should adjust to you. The Back Seat Driver should know about the knowledge and desires of its driver, and act differently because of this knowledge. There are various ways it could act differently.

8.2.1 Improving the directions

If the Back Seat Driver knows what you know about the city, it can give you better directions by using what you know. For example, to a driver who has often traveled up Mt Auburn Street to Watertown Square, and who now wants to get to Belmont Street, it can say "Remember that fork on Mt Auburn, just after the Star Market, where you bear left to get to Watertown Square? Well, Belmont street is the road that is the right branch of that fork." A user who knows about a city no longer needs instructions, she needs information about structure. The object description system TAILOR describes mechanical devices with two different strategies. Novice users hear a process description which emphasizes causal connections, and experts hear structural descriptions. Experts do not need the causal information, they can derive it for themselves[63]. The analogy to route descriptions is that process information is instructions.

Knowledge of the user's knowledge is crucial if the Back Seat Driver is to give orienting information as described below. It would be horrible to have the car give the same lecture about the history of Beacon Hill every time you drove over it.

8.2.2 Learning about the driver

The Back Seat Driver should acquire a model of the user by itself, without asking or having to be told. For properties which change only very slowly (e.g. the user's

login name on the mail computer) it is more acceptable to ask the user for a value: Properties such as color-blindness, or visual or aural acuity do not change quickly. Rapidly changing properties must be learned automatically.

The Back Seat Driver could learn the driver's reaction time by measuring the time between its speech and the driver's operation of the controls. Renault has developed a prototype car which attempts to tell whether the driver is falling asleep by observing the correlation between small motions of the steering wheel and the attentiveness of the driver [84]. This system does not use a generic model of users, but rather learns the correlation for the driver during a registration period.

The Back Seat Driver should learn the style of instruction giving appropriate for the driver. At present, there is not much to learn other than the settings for the parameters in the user model. In future work, the Back Seat Driver should learn not only these parameters, but also acquire rules (e.g., the rules of classification and obviousness) for the user. To learn the driver's preferences for route description requires either observation of the driver herself giving instructions (applying the Golden Rule - "Give instructions unto others as you would have them instruct you") or getting feedback from the driver about the instructions the system provides. One opportunity for learning the driver's style is during the acquisition of speech recognition templates (for user-dependent speech recognition). The new user should play the role of a "back seat driver" and give instructions, while in a car, for some route. The instructions must be given while driving either a real car or a close simulation (such as an "Aspen" movie map) because the form of static driving instructions is much different from real time instructions. Given some a priori knowledge about the ways that a route can be described, it is not impossible that the system could understand the instructions, and infer style from it. A difficulty here is that if the driver knows the route well, many things will seem obvious to her that would not be obvious to another person.

The driver can provide feedback about the suitability of the Back Seat Driver's instructions either explicitly or implicitly. One explicit indication of comprehension is how often the driver hits the "what now?" button. The system might collect long-term statistics on the types of intersections where the user most often requests help, and begin to offer instructions without being asked. Just as the user can ask for more talking with the "what now" button, the Back Seat Driver should provide a "shut up" button (or recognize spoken words to the same effect). The Back Seat Driver could also learn the effectiveness of its directions simply by watching the driver's performance - in particular, her errors. Perhaps it can learn which cues are most effective.

The Back Seat Driver should understand the driver's plans and goals. This is already partially true. When you enter a destination address, you also tell the system that you have the goal of getting to that address. The Back Seat Driver might guess at higher level plans from knowledge about your destination, and take actions to assist you with those plans. To do this, it must know what kind of thing is at your destination address. For instance, if the address you provide is that of a store, the Back Seat Driver can guess that you are going there to purchase something, or at least to do business there. It could then check the hours that the store is open.

8.3. Confidence

People will only use the Back Seat Driver if they trust it. The best way to keep a user's confidence is to never make a mistake. But the Back Seat Driver does make mistakes, and it will make mistakes in the future. What can the system do to regain the user's confidence after an error? It depends on the reason for the error. The most likely causes for error in a future Back Seat Driver are navigation errors and database errors.

When differential GPS is available 24 hours a day, the Back Seat Driver can expect position accuracy to better than two meters. But until that time, there will be navigation errors. The Back Seat Driver will need to model the uncertainty of a position. For instance, when two roads diverge at a narrow angle, it will be unable to distinguish which was taken until some distance passes. It should probably assume that the driver made the correct choice rather than taking the risk of making a false accusation. If it reaches a place where the difference is crucial, yet unknown, it is probably better for it to confess its uncertainty, and say something like "I'm not quite sure where we are, but if you can take a right here, do it, and if not, keep going, and I'll figure things out better in a minute." Or it might ask the driver to pull over and stop (if the driver is at a place where that is safe) to allow for a better position estimate by averaging a few successive estimates. Navigation inaccuracy due to Geometric Dilution of Precision (GDOP) (page 130) is predictable. GDOP from LORAN is function of position alone. GDOP with GPS depends on the positions of the satellites, but position is predictable. The Back Seat Driver could even warn the driver that a planned trip is better deferred because of poor conditions.

Errors in the database will also persist because the database will always be somewhat out of date. The Back Seat Driver can regain at least a little confidence by how it explains the mistake. Suppose that the Back Seat Driver is intending the driver to turn onto "Apple" street, and says "Take a right at the next light". Unbeknownst to it, a new traffic light has been installed at "Pear" Street, so the driver turns there. The Back Seat Driver at present says "I meant for you to go straight." This message is somewhat confusing, because the driver may think that the program has not been consistent. ("First it said turn, so I did, and then it said I should have gone straight.") A better message would be "I did not mean for you to turn onto Pear. I thought that the next set of lights was at Apple Street."

8.4 Speech interface issues

There are two issues concerning speech and the the Back Seat Driver. It would be desirable to make its speech more easily understood, and speech input should also be investigated.

8.4.1 Speech output

Synthetic speech is more difficult to understand than natural speech or digitized speech. Some users complained about speech quality, and others did not understand what they heard. The most common error was to misunderstand the name of a street. This is to be expected, since there are plenty of redundant cues to the meaning of a sentence, but none at all for a name. With practice, people become more able to understand synthetic speech, but even those with a great deal of experience (such as the author) still misunderstand words.

It might be possible to obtain more intelligible speech by using digitized speech rather than synthetic speech. Digitized speech will require a great deal of storage space. There are more than 2000 different street names in Boston. Allowing another 500 words for the actual instructions, and assuming an average duration of .5 seconds for each word, coding this vocabulary at 64 kilobits per second would require 10 megabytes of speech storage. Given that a future Back Seat Driver will use a CDROM for the map, there should be no great difficulty in storing digitized speech on the disk as well. Coded speech would be more intelligible than synthesized speech, and less costly, but not as flexible. It would be impossible to read electronic mail using only stored vocabulary speech.

8.4.2 Speech input

The Back Seat Driver would be much easier to use if you could simply talk to it instead of using the cellular phone keypad. Back Seat Driver could use speech recognition for entering addresses instead of spelling names with a keypad. Address entry is a difficult task for speech recognition for the same reason it is hard for a human to understand machine speech - there are few constraints on a name. No speech recognizer available today can handle a 3000 word vocabulary with acceptable error rates. The car would also have to stop while the driver was speaking, because noise in moving cars for frequencies below 400 Hz can exceed 80dB[62].

Back Seat Driver could also use speech recognition to replace the "what now?" and "what next?" buttons. This is a more tolerant application for speech recognition because there are fewer words to recognize. It may be possible to dispense with word recognition altogether, and simply recognize paraverbals, as in [76].

8.5 Understanding the route

The Back Seat Driver should help drivers to understand the route it gives. This would make the system more pleasant to use. Some drivers complained about being told what to do without any overall presentation of the route. It would also make it easier to follow routes because a driver who understands the route and the city will use that knowledge to help interpret the commands Back Seat Driver gives. For example, a driver following a route that crosses the Longfellow Bridge and who knows that Main Street leads to the bridge knows about how long she'll be traveling before her next instruction, and she'll be attentive to cues at the right time, instead of having to constantly be on the lookout.

The Back Seat Driver already provides a little information about the route. While following a route, a driver can hit a button and hear the total length of the route and the remaining distance. A second button provides the name of the street the car is on (often hard for the driver to know), and the compass direction (almost always a surprise in Boston). But this is just a beginning.

A route should fit into a larger model of the city. This means that the Back Seat Driver itself must have a model of the city and should speak of the route in terms that relate it to the city. There are several opportunities to do this. At the beginning of the route, the driver might hear an overview of the route, naming the major paths followed and neighborhoods crossed. During the route, locations could be described not just as street address but in larger units of neighborhoods and districts. The system might say not, "You're at 900 Mass Ave," but rather "We are now half way between Central and Harvard Squares." Orienting information can be included in instructions, or it might come between instructions, as a passing comment.

Benjamin Kuipers presents a formalization model of people's ability to learn to navigate[46, 47]. In his model there are three stages of spatial knowledge:

- Sensorimotor Procedures: knowledge expressed as conditional procedures: "When you see this, do that"
- Topological Relations: containment, connection, and order
- Metrical Relations: distance, direction

A route expressed in sensorimotor form is a sequence of cues and things to do. This is sufficient to follow a route, but does not support reasoning about the route. You can not reverse the route, because each cue leads only to the next, in the familiar order, and not the preceding. You may not even be able to give the route to someone else, as each cue may be recalled only in the context of the one

EXHIBIT

2

(Part 7)

just previous. This accounts for the familiar "I can take you there, but I can't tell you how." Topological knowledge tells you, for instance, that Massachusetts Avenue enter Cambridge from Boston by crossing the Charles River at the Harvard Bridge, then passes MIT, Central Square, Harvard Square, and Porter Square, and continues into Arlington. Topological knowledge tells you that if you travel west from Somerville you are certain to eventually hit Massachusetts Avenue. It is knowledge of topological structure that gives us the confidence to explore, knowing that we will eventually cross a known feature, so we can always find our way home. Topological knowledge is essential to feeling at home in the city. Metrical knowledge is absolute knowledge of distance and direction. The Back Seat Driver has near-perfect metrical knowledge of the city - it can specify the distance and direction of any point from any other. Metrical knowledge is essential for efficient route finding. People's mental maps of the city usually have correct topology but distortions in metric structure.

Sensorimotor knowledge is the first to be acquired. The instructions the Back Seat Driver gives are sensorimotor instructions - they have to be, since only sensorimotor commands are directly executable. Drivers first learn routes as sequences of sensorimotor instructions², and gradually acquire topological and metrical knowledge. (Kuipers' program is a model of how this learning occurs.)

A better version of the Back Seat Driver should help people acquire topological and perhaps metrical knowledge of the form of the city. People learn from experience, but they can also use other sources of knowledge. People learn metrical relations partly by seeing maps, so it is possible that they can learn topological relations by being told. This is the justification for including orienting information

²I do have a concern that people who depend completely on the system for decisions about which way to turn may never acquire this knowledge, since they would not use their own judgments. In my experience, I learn a route much faster when I'm the driver than when I'm a passenger. Partly this must be because the task of driving forces me to pay more attention to where I am, but also there must be an effect from trying to remember a previous route, or trying to interpret instructions or a map. It would be a pity if the Back Seat Driver actually impaired learning.

in Back Seat Driver. It is also the only reason to show a map. Maps may not be a good way to give directions, but they are an excellent way to show the topological and metrical structure of the city.

8.6 Other directions

The Back Seat Driver helps you get from where you are to some other location. There are some related services that the Back Seat Driver could easily provide.

- It should be able to give the location of a place without giving directions, e.g. the description of Belmont Street in the example above.
- It should be able to give the directions all at once (as does Direction Assistance), and
- It should be able to give directions between any two places.

You probably will not often want these services, but it is easy for the Back Seat Driver to provide them. You might want them because you want to tell someone else how to get to where you are. How frustrating it would be to know that your car could give you the information, but will not.

8.7 Changes to the map

A problem for a practical Back Seat Driver is how to keep the map accurate, since the street network is constantly changing. Over time, new streets are constructed, old streets are renamed or closed. These kinds of changes are predictable, slow, and long-lasting. Other changes are unpredictable, quick, and transient. A

road may be closed for repairs for the day, blocked by a fallen tree, or full of snow. Such changes are usually short lived. Only disasters such as landslides or earthquakes have effects that persist more than a day.

Suppose you are getting in your car for your daily commute. This morning, a road you usually use is closed for repairs. Although you do not know this, the Back Seat Driver does. When you start the car, it can tell you about the closure, and offer to find a different route for you. The most likely means of distributing such changes is to broadcast them on a data subcarrier on FM radio[27]. Such changes might be broadcast once per hour. Your car would be listening even when parked, and could store the changes in local memory.

8.8 Integration with the city

Back Seat Driver should part of a larger system, linking the driver of the car to her home and office, favorite service station, and other cars. When you drive into the city, your car should call ahead to parking lots near the destination, find one with an empty space, and make a reservation for you³. Your car should keep track of its own health and arrange maintenance appointments with your favorite service station. When you are driving home, your car should call home to see whether anything is needed, and if so, give you a route that passes by a store. Research in Japan on the Advanced Mobile Traffic Information Communication System (AMTICS) might lead to such integration. The committee is sponsored by a consortium of the national police agency and manufacturers of automobiles, telecommunications, and electronics equipment, and is developing standards for communications protocols, database formats, and interfaces between human, telecommunication equipment,

³The designers of ERGS anticipated this as early as 1970.

and computers. A European program, Prōmēthēus (PROgram for an European Traffic with Highest Efficiency and Unprecedented Safety), has similar aims.

When cars have positioning systems, it will become useful to the public if these cars can transmit their positions and routes on request. If enough cars report their positions and speeds, then a central agency can estimate traffic flow. This information will be useful to route planners in the cars. Furthermore, the agency can attempt to improve the traffic flow by balancing load. This could save a lot of money. Widely available positions and routes would also be useful in giving directions. If the Back Seat Driver knew that the car just ahead of you was about to make a turn it wanted you to make (by asking the navigation assistant on that car) then it could give you directions by just saying "Follow that car".

8.9 Policy

Before manufacturing a Back Seat Driver, there are public policy issues that must be addressed. There is a question of liability, and also some issues of policy on privacy. I do not have any answers to these issues, but it is important that they be raised.

8.9.1 Liability

Sometimes the Back Seat Driver gives orders that should not be obeyed. It might make a mistake, and send the driver down a one-way street, or it might just tell the driver to "keep going" when traffic is in the way. Drivers who become accustomed to obeying the voice without thinking for themselves will have accidents.

The most likely source of error is the map database, which can easily be inaccurate. Manufacturers of maps are subject to contract liability law and negligence law. The latter is the greater threat because of the greater potential size of damage awards. According to [26]:

The manufacturer/seller may be held responsible for any injury to the buyer, or to any other user, or to any foreseeable bystander while the device is being used. The injuries may include direct physical or economic injuries as well as indirect, intangible injuries. In addition, there is the possibility of punitive damages as a penalty for carelessness.

Furthermore "there is little that a seller can do to protect himself from frivolous negligence suits". Once maps (or systems using them) are mass produced, they become subject to product liability law, which has the same potential for damages, and a lesser burden of proof:

The buyer/user only needs to prove that he was injured, and that the defect caused his injury. The manufacturer of the product becomes liable to the same extent as if direct negligence had been proved....there are very few defenses that are successful in strict liability cases.

Rogoff[69] proposes that the federal government issue electronic maps on the grounds that the government can not be sued for liability.

8.9.2 Privacy

A second issue is the question of privacy. The Back Seat Driver knows if you drive too fast or fail to stop at a stop sign. What should it do? Should it automatically slow the car down, sound an alarm, or just record the violation

silently, and later tell the insurance agency - or the police? Will automotive companies seek to record engine statistics to verify that the driver treated the car gently during the "break in" period?

The Back Seat Driver might remember every place it has ever been and the time it was there - there are good reasons for it to do so - but this record might be something the driver would not want anyone else to know about. Sometimes drivers convicted of drunken driving are allowed to keep their licenses, but only for driving to and from work. It is easy to imagine a demand for using a position tracking system to verify compliance with this restriction, since some jurisdictions already use personal locators for similar purposes. Regulations are needed to ensure that the trip history in your car is your private data, not to be inspected without a search warrant, and to require that all vehicles be able to take trips "off the record". These issues become even worse if cars broadcast their positions or routes, as proposed above. I am not at all sure how to resolve the conflict between privacy and utility. I leave it as a problem for the reader.

Chapter 9

Summary

The Back Seat Driver is a working prototype of an intelligent automotive navigation assistant. The important differences between the Back Seat Driver and other such programs are that the Back Seat Driver

- finds routes for the driver, instead of simply displaying position on a map,
- tells the driver how to follow the route, step by step, instead of just showing her the route, and
- speaks its instructions, instead of displaying them.

Each of these design goals has required new features in the Back Seat Driver or its street map database:

Finding legal routes requires that the street map database include legal connectivity in addition to physical connectivity. Finding routes that are fast requires knowledge of expected travel rate, which includes quality of the street and the presence of traffic control devices such as traffic lights and stop signs.

Giving instructions for route following requires breaking the route down into a sequence of driving acts: it requires a taxonomy of these acts, and knowing when

an act is obvious to the driver and when it needs to be mentioned. This further requires knowledge about the individual driver, for what is obvious to one may not be so to another.

Speaking instructions requires generating English language descriptions of actions. These instructions tell the driver what to do and when to do it. Instructions must be expressed in terms familiar to drivers – in particular, they must use timing and landmarks, rather than distance, to tell drivers where to act. This in turn requires that landmarks be present in the map database. The use of natural language, in contrast to displays, affords the ability to refer to past, present, and future places and events – thus the Back Seat Driver may refer to past mistakes, current actions, and future landmarks with equal ease. It would be more difficult to convey this information by graphical means, moreover it would require that the driver look away from the road to gain navigation information.

Speech, especially synthetic speech, as an output medium imposes constraints on the interface. First, the transient nature of speech requires that the program be prepared to repeat its utterances on demand, which means the ability to construct a new utterance with the same intent, not necessarily the same words, as a previous message. Synthetic speech being sometimes hard to understand, the program must choose its words to provide redundancy in its utterances.

Although the Back Seat Driver is working now, much work remains before we can expect to see a system like it used by ordinary people, whether they are tourists, commuters, or taxi drivers. Automakers remember the public rejection of earlier “talking cars”, which spoke only of unlocked doors and empty fuel tanks. Nobody wants a nag in the car, and therefore a Back Seat Driver must know about the knowledge, desires, and goals of its driver, so that it will know when to speak, and when to be silent. This knowledge will be difficult to acquire, but if it can be done, our culture may come to have more favorable meaning for the phrase “Back Seat Driver”.

Appendix A

Location Systems

A navigation assistant such as the Back Seat Driver has to know the position of the vehicle. In this section I present a review of the available technology for determining location. To motivate this, I begin with a discussion of the requirements of the Back Seat Driver.

A.1 Accuracy requirements of the Back Seat Driver

At the minimum, a direction giving system must determine its position to the nearest block. If it is to tell the driver when to turn it must be able to distinguish between the closest two adjacent turns. Hall estimates that a 50 meter accuracy is sufficient because he takes the average length of a block to be 100 meters[32]. Pilsak demands 25 meter accuracy and 10 degrees of heading resolution to select among narrowly diverging intersections[65].

Consideration of the Boston street map shows that it has many streets which

are both short and a possible choice point. Figure A-1 shows the percentage of segments which are shorter than a given length. Only segments with at least one choice point are counted here. If the Back Seat Driver has only 25 meter accuracy

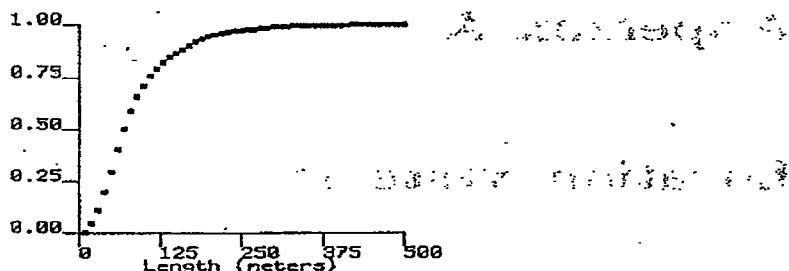


Figure A-1: Fraction of segments shorter than X in meters

then it will not resolve more than one fourth of all streets. An accuracy of 10 meters is more desirable.

The Back Seat Driver's use of speech imposes more strict requirements on position because of limitations on time. Unlike a display, speech is transient. An action described too soon may be forgotten. The Back Seat Driver is intended to speak at the latest time that still permits the driver to act. Allowing two seconds for speech, a car at 30 mph covers 27 meters. This suggests a minimum accuracy of 15 meters.

A.2 Classifying location systems

Following Cooke[15], location systems can be divided into two categories:

- Position finding systems determine position directly by detecting an external signal.

- Position keeping (or dead reckoning) systems estimate the current position from knowledge of an earlier position and the change in position since that position.

A.3 Position finding

All existing position finding systems use radio signals. The broadcast stations may be located on street corners, seacoasts, or in orbit around the earth. Systems differ in their range, accuracy, and cost. A survey of positioning systems organized by operating frequency appears in [74]. That survey includes many specialized, limited range systems. The list here includes only those systems which might plausibly be used in for automobile navigation.

A.3.1 LORAN-C

LORAN (LOng RAnge Navigation) is a navigation aid originally intended for ships, now extended to civil aviation and land navigation. It relies upon measuring the difference in arrival time between signals transmitted simultaneously from three separate known locations. If the receiver is equally far from two transmitters, the signals will arrive at the same time. If the receiver is closer to one, that signal will arrive first.¹ A measurement of the inter-arrival time for a signal from two positions fixes position to a hyperbolic line, called a line of position (LOP). The line of position takes the shape of a hyperbola because the locus of points with a constant difference in distance to two given points is a hyperbola. The estimated position is the intersection of two lines of position.

¹This is a simplification. LORAN transmitters do not really transmit simultaneously, the receiver compensates for this.

The accuracy of a LORAN position fix depends upon three categories of factors: geometry, radio propagation, and equipment. The significance of geometry is known as Geometric Dilution of Precision (GDOP) and is illustrated in figure A-2.

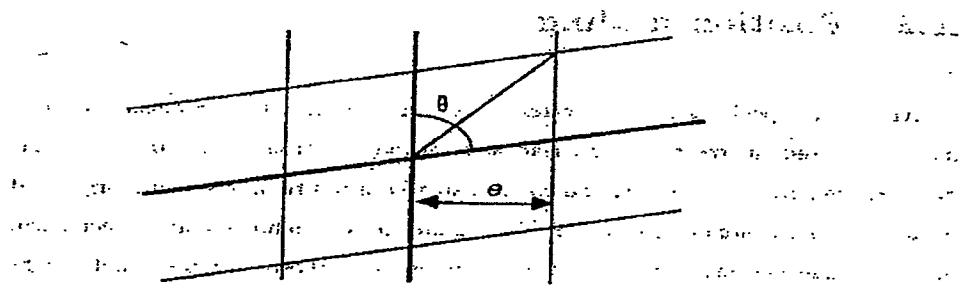


Figure A-2: Geometric Dilution of Precision

The placement of the LOP is subject to some error. If the maximum size of the error is e , then the true LOP will be parallel to the measured LOP, and at most a distance e from it. The two LOPs might meet at any angle, depending where the receiver is. If the LOPs meet at right angles, errors in each are independent, but otherwise the errors add. The maximum error, in the case where the error in each LOP is e is the length of the base of the isosceles triangle whose sides are e . This is $2\cos(\theta/2)$, where θ is the angle between the two lines of position. So when the LOPs meet at right angles, this is $\sqrt{2}e$.

Radio propagation affects LORAN by causing erroneous time differences, thus spurious position errors. Factors affecting radio propagation include weather, season, time of day, and signal path. Radio propagates more slowly over land than over water, and is reflected or delayed by topography and man-made features[10].

Equipment conditions at transmitter and receiver can also introduce false time differences. Dutton estimates a cumulative error of 15 to 90 meters when within 200 miles of the transmitting station.

Accuracy can be improved by means of differential LORAN[10]. A differential LORAN system requires setting up monitoring stations at fixed, known positions. These stations compare the LORAN estimated position to the known position to obtain instantaneous corrections to the radio signal. Mobile LORAN receivers can use these corrections over radius of approximately 35 miles. Blizzard estimates that differential LORAN can provide repeatable accuracy to ± 20 meters. Cooke cites an accuracy of 5 meters.

LORAN is used in at least two commercial vehicle location systems, the Vehicle Tracking System sold by IImorrow, Inc, and the Automatic Vehicle Location system of Motorola[73]. IImorrow claims an accuracy of 22 to 36 meters for their system. Janc (the patent holder for the Motorola system) provides an account of sources and the worst case magnitudes of errors for LORAN, and measures a 500 meter accuracy with 500 meters with a 95% confidence interval [39]. El-Sawy reports on a LORAN vehicle tracking system in [23]. He does not give accuracy figures, but says "Some system operators have voiced concern over the inconsistency of the vehicle tracking in some geographic areas". In rural areas without a complicated signal path, McGillem obtained accuracies of 45 meters[55].

A problem that remains is that LORAN does not completely cover the continental United States. The system is being expanded on behalf of civil aviation users and should be complete by 1990[78].

A.3.2 Polled pulse time ranging

In LORAN, vehicles determine position by comparing arrival time difference of signals from three (or more) transmitters. An alternative is for three or more

fixed receivers to compare the arrival time of a signal sent from the vehicle. In 1977 Hazeltine developed a system which does just this. A central site broadcasts a prompt, and transponders on vehicles echo the prompt. The obvious drawback is that there must be a way to distinguish each vehicle's response. The Hazeltine system uses time multiplexing. The vehicles do not respond all at once, but rather one at a time. Each vehicle has a time slot for its response which lasts long enough that any previous response will have been received by the time the vehicle transmits. In the Hazeltine system, time slots are two milliseconds long, so 30,000 vehicles can be tracked in each minute. The system was tested in Dallas[66] and Philadelphia[11] and achieved accuracy of 300 feet within a 95% confidence.

A.3.3. GPS.

The Global Positioning System (also called NAVSTAR) derives position by measuring the travel time of (and thus distance to) radio signals from satellites orbiting the earth. The system is incomplete at present, but will eventually consist of 24 satellites. A GPS receiver can determine both position and velocity in three dimensions, providing that four satellites are visible at the time. If elevation is known (or independently measured) then only three satellites are required.

GPS provides two classes of service. The Precise Positioning Service (PPS) offers accuracies to eight meters, but is only available to military users. The Standard Positioning Service available to the public provides only 200-meter accuracy. One satellite can provide two levels of service by broadcasting position and time information at two different levels of accuracy, using a different code for each. SPS information is encoded with the Coarse Access (CA) code which has been made public. PPS information is available only to those who can decode the P code.

GPS accuracy is influenced by the same factors that influence LORAN: geometry, propagation, and receiver, and by others. The geometric effects are more

complicated because GPS satellites move. The GDOP for a LORAN fix depends only on one's position on the Earth, but the GDOP for GPS changes over time. The problem of unpredictable propagation is also worse because the signals travel through several different regions of the atmosphere. There is a well accepted model for ionospheric propagation. The GPS signal includes parameters for this model, enabling the receiver to compensate for the effects. The effects of the troposphere are less well understood. GPS satellites transmit on two different frequencies, L1 (1575.42 MHz) and L2 (1227.6 MHz). The L1 channel carries both C/A and P code, the L2 carries P only. PPS receivers can compensate for the effect of the troposphere by comparing the effect on L1 and L2, since this effect depends upon radio frequency. GPS accuracy also requires that the satellite orbits be predictable and stable. Under normal conditions, this is so, but intense magnetic "storms" can raise the outer atmosphere sufficiently that satellites feel a noticeable drag. The geomagnetic storm of 13 March 1989 caused so much drag that GPS tracking was unavailable[34].

As with LORAN, differential GPS is possible[40, 9]. It should be possible for a GPS receiver at a fixed, surveyed location within the city to measure the errors from propagation delay and broadcast correction factors, perhaps over single sideband FM, to receivers in cars. Experiments with differential GPS have shown accuracies of two to three meters using the C/A channel[45, 2, 3]. At least one expert claims that sub-meter accuracy is possible[7]. Surveyors using so called kinematic positioning can obtain even better accuracy but require tens of minutes to do so. It is unclear whether this accuracy can be obtained in real time.

A.3.4 TRANSIT

The Navy Navigation Satellite System system, also called TRANSIT, is the older of two satellite navigation systems. NAVSAT differs from GPS in that only

a single satellite is required for a position measurement, instead of three or four for GPS. This is possible because NAVSAT satellites travel in a much lower orbit (approximately 600 miles) than GPS, and thus move faster, covering a larger distance during a fixed time. A single NAVSAT position measurement requires about 10 minutes. In this time, the satellite moves more than 2000 miles thus providing a large baseline for calculation. Equipment at the receiving station measures the Doppler shift in the received signal to obtain a measure of the satellite's velocity relative to the ship. From this, and the knowledge of the position of the satellite, the receiver's position may be measured. Under ideal conditions, a NAVSAT position fix is good to within about 45 meters.

NAVSAT alone is not sufficient for car navigation because of the long time to obtain a position measure and the long time between satellite orbits. Even under ideal circumstance, one can obtain but one position measurement every 90 minutes, but at times the period between useful satellite passes can be as high as 16 hours ([80]p 190). The vehicle's velocity must be carefully measured while obtaining a position fix. An error of as little as one mile per hour can cause a position error of up to 1/4 mile in the position obtained. NAVSAT might serve to supplement a position keeping system. It is unlikely to do so, because it is being shut down in 1994.

A.3.5 Other satellite systems

The Soviet Union's satellite positioning system, GLONASS, is quite similar to GPS[16]. Like GPS, it consists of a highly accurate, encrypted signal and a less accurate clear signal. Details on the clear signal have been released and tests in the West show accuracy similar to that of GPS[17].

The STARFIX system[20, 61] operates in a manner similar to GPS, but uses geosynchronous communication satellites rather than dedicated navigation satellites. A ground station sends an "event" message to the satellites, which relay it to STARFIX users and to reference stations. The reference stations broadcast a local differential correction to users in their area. Tests have shown an accuracy as good as 2.5 meters, degrading to 30 meters at maximum range from a reference station (1000 miles). STARFIX appears to have suitable accuracy for direction giving, but present receivers are much too large (ten cubic feet) and heavy (300 pounds) for cars.

GEOSTAR[67] is a proposed, but not yet operating, commercial system to provide two dimensional position and two way digital communication to users in the continental United States. Unlike GPS, GEOSTAR users are not passive. The system determines location by broadcasting an interrogation to a user, which then transmits a response to at least two satellites. The satellites relay the response to centralized processing facilities on the ground, which determine the user's position by measuring the elapsed time for the signals. A network of benchmark receiving stations provide increased accuracy by differential techniques. Position accuracy is said to be within seven meters. Note that the mobile station does not know its position unless the central facility transmits it. Note also that since only two satellites are used, the user must supply the altitude.

The European Space Agency is studying an independent satellite positioning system, also called NAVSAT, which would provide 12 meter accuracy to its users.

A.3.6 Beacons

A less exotic alternative is to place radio (or infra-red) beacons or "signposts" at intersections. These provide position, or rather, an indication of proximity. Such a system was installed in Huntington Beach, California[29, 31]. In this installation,

482 transmitters were installed on utility poles approximately every quarter mile. Huntington Beach is ideal for such a system because it is laid out in a grid. The system provides position to 300 feet with 90% confidence on the major streets, and 600 foot accuracy on the lesser streets. The disadvantage of such systems is that they require a large investment in installing transmitters at each intersection.

Signposts can also be passive, and interrogated by the vehicle. Proposed systems include corner cube reflectors tracked by lasers[86], magnetic strips embedded in the road[48], and transponders which respond only when interrogated[41].

A.3.7 The cellular phone system might serve as a location system,

I speculate that one might be able to use the network of cellular telephone base stations for a positioning system. This is possible because cellular phone sites already broadcast a status message on a regular basis. This is how cellular phones determine which frequencies are available and which carriers (telephone companies) are in operation. All that is required is that cellular phone sites be synchronized to a common clock. The positions of cellular base stations are well known (and can be known in the car), so it is a simple matter to measure distance to the cellular phone stations. In effect, this would make cellular site transmitters into earthbound GPS satellites. An alternative is to use polled pulse time ranging. Here, the cellular site would compute the position of the car and retransmit it².

²I do not know how to estimate the theoretical position resolution. The carrier is around 800 MHz. I do not know what the envelope of the carrier looks like when turned on, or how accurately the arrival time of the waveform could be measured. If measured to the period of the carrier, that is 1.25×10^{-9} seconds, or about 40 meters, which is good, but not good enough. But could you measure phase and do better?

A.3.8 Vision

The defense department is sponsoring research on an Autonomous Land Vehicle. This vehicle will use machine vision to determine location. The research is far from ready for application, but we can anticipate a day when navigation systems will determine where they are by just looking out the window, in the same way that humans do.

A.4 Position keeping

Position keeping, also known as *dead reckoning*, obtains position indirectly, by keeping track of the displacement from an originally known position. One can measure displacement directly, or they measure velocity or acceleration and integrate over time to obtain displacement.

The forward motion of a car is measured by the odometer. On late model cars, the odometer cable has been standardized. It revolves once every 1.56 meters. This is a reliable measure of forward progress, as long as the wheels do not slip. Measuring direction, though, is more difficult. Among the possibilities are:

- **magnetic compass** A magnetic compass has the advantages of small size and ease of use[64], but is unreliable because of variation between magnetic and true north and deviation due to the ferrous material of the car and magnetic flux arising from electric currents within the car. For a discussion of the difficulties of magnetism, see [87].
- **steering wheel** The steering wheel could be instrumented to measure the amount of turning.

- **differential odometer** When a car turns, the two rear wheels travel different distances, and thus rotate at different rates. Measuring the difference in rotation should provide an indication of amount of turning. By lucky coincidence, this differential rate of rotation is just what is measured by anti-skid brakes, so no additional instrumentation is required to obtain this measure. There is no obvious reason why it should not be available on all new cars in the near future.
- **gyroscopic** Gyroscopes measure angular acceleration. The familiar rotational gyroscope and esoteric laser ring gyroscope are not suitable for automotive use because they are too expensive. Lower cost alternatives are the rate gyro and the gas jet gyro. The rate gyro measures angular acceleration in a vibrating piezo-electric substance[43, 25]. The gas gyro measures turn (or yaw) rate. In this design a jet of gas travels down the center of a sealed tube. Anemometers are placed on either side of stream. When the car turns, the stream is deflected and the velocity is measured. The velocity of the gas at the anemometer is proportional to the turn rate. Such devices can measure turn rates of up to 100 degrees per second, with a noise of about one half degree/second.

A.4.1 Dead reckoning systems need correction

A position keeping system with no error could be calibrated when installed, and then maintain its own position indefinitely. Unfortunately, errors arise in measuring both distance and heading. Sources for error include difference in tire pressure, composition and wear, slipping, cross steering from winds, change in tire contact patch in turns, magnetic anomalies, and gyro noise. The NEC dead reckoning system, for instance, accumulates about one meter of error for every ten

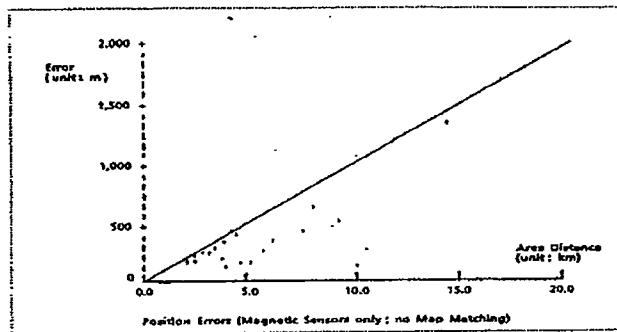


Figure A-3: Dead reckoning errors increase with distance

meters traveled (Figure A-3). The error is even worse when traveling near railroads because the NEC system uses a magnetic compass.

Some dead reckoning systems recalibrate themselves to eliminate systematic errors. Such recalibration is possible when the vehicle is at a known position or when stopped. For instance, if the ETAK navigator finds a consistent overestimation of distance it compensates for tire wear and it adjusts its magnetic compass while traveling on a road with known heading. The Honda Gyroocator adjusts its yaw rate sensor when stopped. Recalibration can also be applied to position finding systems. For instance, a car stopped at a known location could accumulate differential GPS or LORAN corrections on its own behalf.

One way to correct dead reckoning errors is to use knowledge of the road network as a *constraint* on position, in what is known as *map matching*. Map matching requires that the position keeping system have a map of the locale where the vehicle is being driven, and is based on the assumption that the vehicle is always on a street present in the map. If the estimated position falls to one side of the road, the estimate can be corrected. When the vehicle makes a turn, the system assumes the vehicle is at the closest intersection, and thus the absolute position can be estimated. Figure A-4 illustrates the effect of map matching.

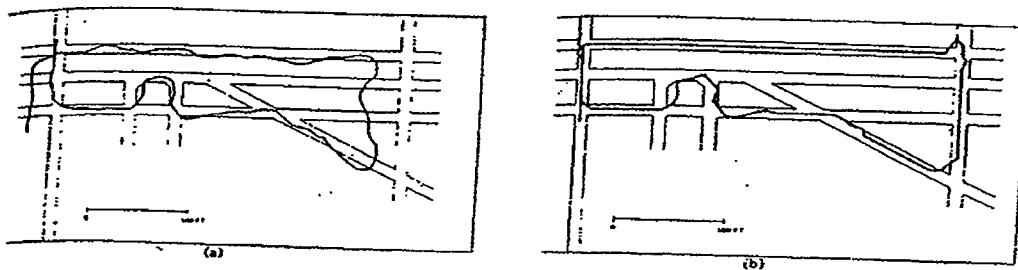


Figure A-4: Map matching corrects dead reckoning errors

Every dead reckoning system uses some form of map matching. Map matching reduces the stringency of position keeping, but accuracy remains a concern, since the system must maintain its position when the driver drives off the map, e.g. into a driveway or a parking lot.

A.5 Hybrid systems

Positioning systems must combine dead reckoning and position sensing, since neither alone will work all the time. A position keeping system needs periodic corrections, but an absolute positioning system that depends on radio reception will not work in tunnels or bridges. Such hybrid systems are discussed in [59, 37, 43, 44, 38]. Karimi stresses that hybrid systems should not simply combine all knowledge sources, but should reason about the expected reliability of each one using knowledge of current road conditions, weather, location and so on. For example, when entering a tunnel, the system should no longer trust satellite positioning data, and it should also allow sufficient time after emerging to reacquire

position. Itoh and others [38] shows the effect of construction and topography on GPS reception. Under good conditions (in open areas), they obtain positions accurate to within 30 meters, but when driving in urban or mountainous areas they find a GPS reception rate as low as 13% (contrasted with 75% when near the ocean). Although this may improve when the full GPS constellation is operating, it seems likely that hybrid systems will remain essential.

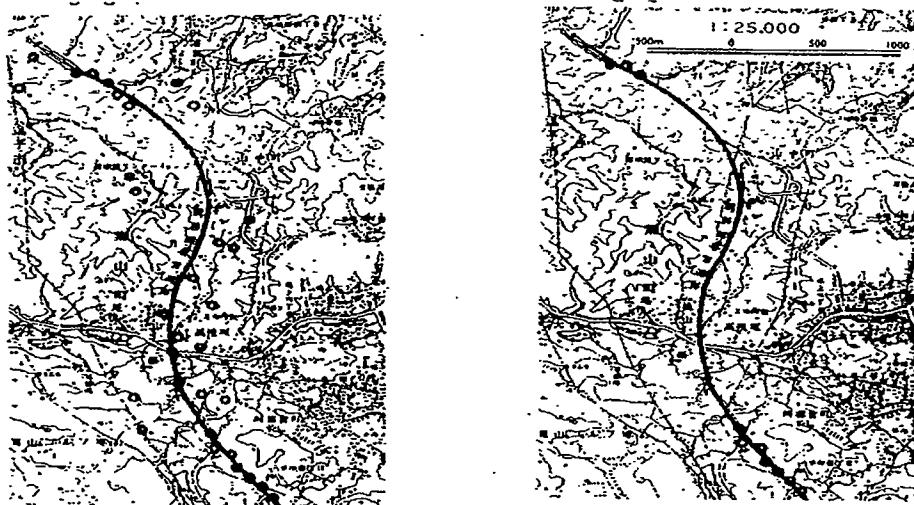


Figure A-5: Dead Reckoning corrects for loss of GPS signal

Appendix B

Route Finding

Finding a route between two points in a street network is an example – perhaps the best example – of searching a general graph¹. The task is to find a sequence of segments that lead from the origin to the destination. There are usually a great many distinct ways of getting from one place in the city to another, some better than others. Ideally, we want the *best* route, not just any route. Graph search algorithms differ in the quality of the solution they find and the time they require. For the Back Seat Driver we want an algorithm that finds a good route, though not necessarily the best, and does so in a short time.

The Back Seat Driver uses an A* search algorithm[35]. The A* algorithm is a form of best-first search, which itself is a form of breadth-first search. This chapter begins with a discussion of breadth-first search, then shows how A* search is built from it.

¹See page 99 for more on graph search.

B.1 Breadth-first search considers all possible partial solutions in parallel

Breadth-first search is named so because of the view of search as exploring a tree of all possible decisions. The root of the tree is the current state. Leading from this route are one or more arcs, each corresponding to a possible action. Each action brings about a new situation from which other actions are possible. A solution is a set of actions leading from the initial state to the desired situation. Search algorithms differ in the order in which they consider the possible actions. In a breadth-first search, the tree is divided into levels, where the first level actions are those leading from the root, the second level actions are those that come from situations after first level actions, and so on. All actions at a given level are considered before any at the next higher level. In this way, the search systematically covers the entire tree, skipping nothing.

In the case of route finding, the root state is the initial position, and arcs correspond to street segments leading away from the initial position. It is important to distinguish the search space, which is a tree, from the problem space, which is a general graph. Even when two paths on the map lead to the same place, they are distinct paths in the tree because they have different histories. After driving all the way around a block, you are in the same place as before on the map, but the search tree has become four levels deeper. The position is the same, but the situation is different. Time and distance have both changed.

While the breadth-first search is operating it maintains a list of all possible partial routes. (A partial route is a sequence of segments leading from the origin to some intersection.) The search procedure is a loop. Each time around the loop, the procedure considers each partial path: For each path, it considers each segment leading away from the intersection at the end of the path. If there are no segments

leading from an intersection, then that path is a dead end, and can be dropped from the list of candidates. If the segment contains the destination, the search is complete. Otherwise, if the segment is not already present on the path somewhere, a new partial path is formed by appending the segment to the end of the path, and this partial path is collected into a second list. (Actually, the segment may be on the path twice, once for each direction of travel. The route finder does not make U Turns in the middle of the street.) After each possible path has been considered, if none have yet led to the goal, the list collected during the loop becomes the new list of possible routes. In this way the search systematically examines every possible path.

Figure B-1 shows an example of breadth-first search. At the beginning of the search, there is just one possible route, leading from the origin to the end of the current street segment. (This corresponds to driving to the end of the street, to the first intersection, in the direction the car is already facing. The Back Seat Driver never asks the driver to make a U Turn in the middle of a street, unless the street is either a dead-end or the driver is somehow facing the wrong way on a one-way street, because such moves are difficult and often illegal.) Here is the first choice point. The route can continue left or right. Neither of these segments is the destination, so there are now two possible routes to consider. In the next step, one possible path generates three more possible paths, and other just one (because the intersection is with a railroad, the only possible next path is to continue straight). Now there are four candidate paths. As the search continues, the number of possible paths grows quickly. Table B.1 shows how the number of potential routes increases at each step. After 1 minute, there are 11,257 possible paths under investigation, but only 1019 segments have been examined. Clearly most of the paths are overlapping. Breadth-first search is wasting too much effort by considering every possible path.

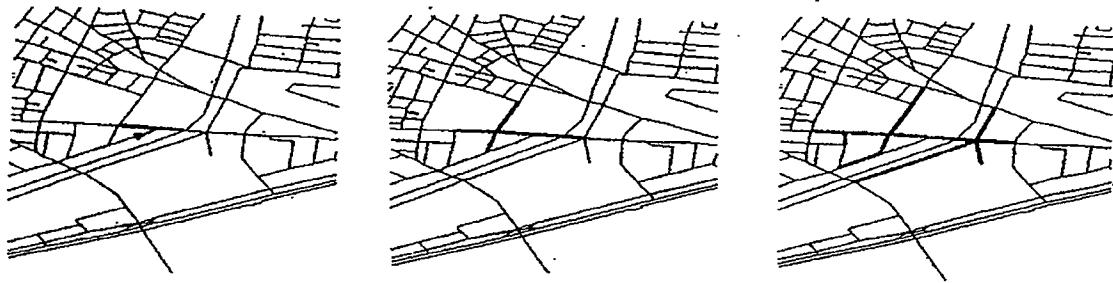


Figure B-1: Breadth-first search example

B.2 Best-first search saves effort

The search procedure given above finds the path with the fewest segments. This is not necessarily the *best* path. Suppose one takes "best" to mean shortest. It is possible (even likely) that there is a path with more segments which is also shorter. To be sure of finding the best path, the search can not stop when the first path is found, but must continue, expanding each path, until *all* paths are complete. Only then can the best be found, by comparing each complete path. This is not at all desirable, since there could be (and in fact will be) many paths. The difficulty arises because the algorithm considers paths in order of their size, not their value. The best-first algorithm solves this problem by keeping track of the (partial) cost of each path, and examining the one with the smallest cost so far. This requires a function that can compare two routes and produce a numeric rating. Such a function is called a **metric**. To further reduce the cost of searching, before adding a segment to a path, the best-first search checks to see whether it is a member of *any* other path. If it is, it is not added, for presence on the other path means that

Step	Number of potential routes
1	1
2	2
3	4
4	8
5	19
7	33
7	59
8	117
9	220
10	397
11	697
12	1199
13	2085
14	3700
15	6469
16	11257

Table B.1: Breadth-first search considers many potential routes

that other path was a less expensive way of reaching the same segment. The other path must have been cheaper, since it was discovered sooner.

B.3 A* search avoids falsely promising paths

Best-first search finds the best solution and requires less time than exhaustive breadth-first search, but it can take a long time because it must consider partial solutions with an initial low cost which prove expensive when complete. The A* algorithm avoids wasting time on such falsely promising solutions by including an estimate for the completed cost when selecting the next partial solution to work on. The cost estimate function is $f^*(r) = g^*(r) + h^*(r)$, where r is a route, $g^*(r)$ is the known cost of the partial route, and $h^*(r)$ is the estimate of the cost to go from the endpoint of the route to the goal. The h^* function must have the property of being always be non-negative and never over-estimate the remaining

cost. An h^* meeting these two conditions is said to be *admissible*. It should be obvious that if h^* is chosen to be always zero, then A* search is just best-first search. In applying A* to finding routes on a map, h^* is just the cartesian distance between the endpoint of the partial route and the destination point. It is certain that no route will be shorter than the straight line, so this estimate is never an over estimate. Table B.2 shows that A* search is more efficient than best-first. The two algorithms were compared on thirty different routes. The routes are identical, but A* took much less time, because it examined fewer potential paths.

	segments touched	time
best-first	355711	2435
A*	59803	455

Table B.2: A* search touches fewer segments, and is therefore faster

B.3.1 A suboptimal, but faster algorithm, is desirable

The A* algorithm finds the optimum route, but the Back Seat Driver might be better served with an algorithm that finds a reasonable route in less time. This is especially true when the vehicle is in motion. The longer the route finder takes, the greater the distance that must be reserved for route finding. As this distance becomes larger, it becomes harder to predict the future position of the car. We can do this by choosing an h^* which multiplies the estimated distance remaining by a constant D . Setting D greater than one makes h^* no longer admissible, since the estimate might exceed the actual cost by a factor of D . The resulting routes are no longer optimal, but are still pretty good. The effect is to make the algorithm reluctant to consider routes which initially lead away from the goal. Table B.3 shows the cumulative path length and search time for thirty different routes with four different values of D . As D increases, the path length increases, but the time

EXHIBIT

2

(Part 8)

decreases more rapidly. In this table *Length ratio* is the ratio of the path length to the length of the optimum path, *Time ratio* is the ratio of search times, and *Payoff* is the ratio of the change in length to the change in time.

Weight	Length	Time	Length ratio	Time ratio	Payoff
1	109.37	455	1.00	1.00	1.00
2	116.81	61	1.07	0.13	7.97
4	122.65	50	1.12	0.11	10.20
8	126.67	45	1.16	0.10	11.71

Table B.3: Comparative search times and route lengths for 30 routes with different values of distance weight factor

The route finder uses a value of 2 for D . This yields the greatest increase in payoff. A possible improvement is to run the route finder twice, first with a high value of D to find an initial route in order to begin the trip, and then with a low D to search for a better route, using spare time while driving.

Appendix C

Communication with the car

The Back Seat Driver is a prototype of an in-car navigation system, but it was actually implemented on a large workstation computer¹. This computer is too large to fit into the car, so instead I used cellular phones to carry data from the car to the computer, and voice from the computer to the driver. This chapter describes the actual experimental setup. It is of little theoretical interest, but may be of practical value to others attempting to send data through cellular phone links, in this area.

The workstation communicates with the driver and the onboard hardware through cellular phones, as shown in figure C-1.

The position sensor estimates vehicle position, heading, and velocity, and sends a data packet, once per second, through the modem to the workstation. The workstation sends characters to the Dectalk speech synthesizer, which in turn sends voice over a second phone to the driver.

¹a Symbolics Lisp Machine

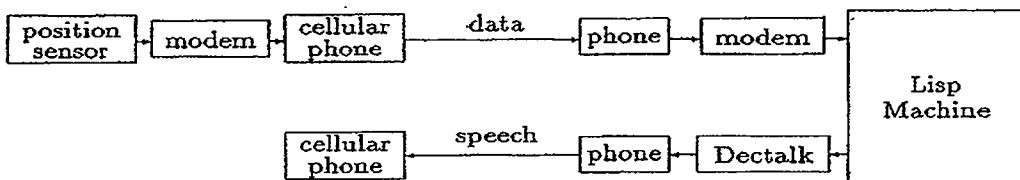


Figure C-1: Communications block diagram

C.0.2 Cellular phones are hostile for data transmission

Nearly everyone who has used a cellular phone knows how noisy they are. Cross talk is common. On several occasions I have heard one and even two other conversations at the same time. Noise bursts and signal loss make it hard to hear. A sufficiently bad noise burst can even cause the cellular system to terminate the call. The problems for data transmission are even worse[6]. By its very nature, cellular-radio transmission is subject to multi-path interference, which causes periodic fades as the antenna moves in and out of anti-nodes. In addition to this fading, there is a complete loss of audio signal for as long as .9 seconds when the phone switches from one cell site to another (hand off).

My attempt to use an ordinary (land-line) modem from the car² was unsuccessful. Even at 300 baud the connection was too noisy to use. Worse, connections seldom lasted more than five minutes. In all cases, I set the "loss of carrier" register (S10) to its maximum value, 20 seconds. Loss of carrier signal alone can not explain why the connections dropped. The modems were capable of tolerating a complete loss of audio for up to twenty seconds.

²I used a Worldlink 1200 from Touchbase systems in the car, with a Morrison and Dempsey AB1 data adapter, and an NEC P9100 phone, boosted to 3 watts. At the base station I used both a Practical Peripherals 2400 and a Hayes Smartmodem 1200.

I had better results using an error correcting modem³ made by the Spectrum Cellular Corporation. This modem uses a proprietary protocol (SPCL[60]) for error correction. The Spectrum product virtually eliminated noise, at the price of a lower data transmission rate. According to the protocol, the transmitting modem groups characters into packets that include error correction bytes. If only a few errors occur, the receiving modem repairs the data and acknowledges receipt. If there are many errors, the packet is retransmitted. If the sending modem has to retransmit too often it makes the packets smaller, on the assumption that a smaller packet has a better chance of success. This is less efficient, since packets have a fixed overhead, the percent of the channel used by data decreases. When conditions improve the modem increases packet size again. In theory, the modem can transmit at 120 characters per second, but I estimate an average value closer to 30 characters per second. I made this estimate by recording the time required to receive the three characters of an odometer sequence. This sequence is transmitted once per second. Figure C-2 shows a histogram of durations for the three character sequence. The mean for this histogram is 94 milliseconds, which is 31 milliseconds

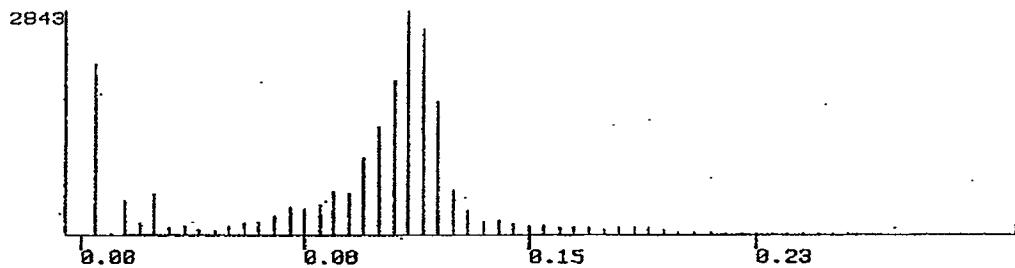


Figure C-2: Histogram of durations of odometer sequence

per character, or 32 characters per second. Tests by Fontana[24] found better results, between 75 to 80 characters per second.

³The "Bridge"

Even with the cellular modem, calls are sometimes dropped. Figure C-3 shows the probability of a call being dropped plotted against time. The measure of probability is obtained by comparing the number of calls dropped at or before time t with the number of calls that lasted at least that long. The call durations are usually long enough for a successful trip with the Back Seat Driver. Voice calls are dropped at about the same rate as data calls.

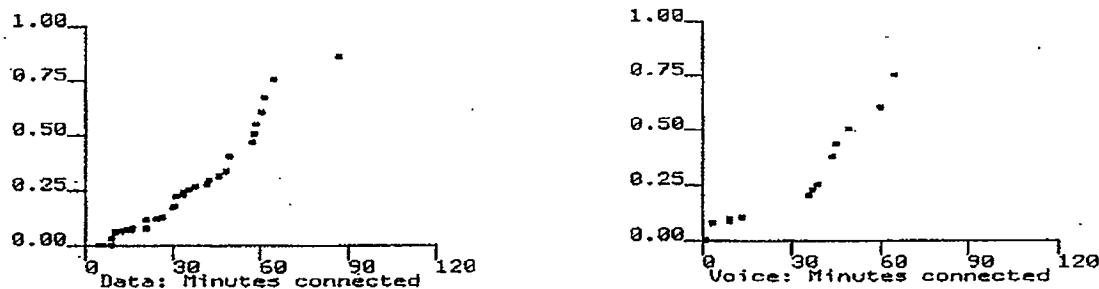


Figure C-3: Probability of cellular call termination increases with time.

C.0.3 Retransmission introduces latency

The protocol used by the Spectrum modem acknowledges all data transmitted. If the acknowledgment is not received, it retransmits the data until acknowledged. Under adverse conditions, this can result in an arbitrarily long delay. This is a problem when real-time data is transmitted. Observation of the Back Seat Driver shows that sometimes the system will "freeze" for from one to ten seconds. During this time, the car of course continues to move. If these freezes occur near decision points, the driver may go past the intersection without being told what to do. At

20 miles per hour a car travels nearly 45 meters in five seconds. Figure C-4 shows a closeup histogram of the average arrival rate of odometer packets. The navigation system in the car sends a packet once every second. Most packets arrive within a second, but a few are delayed, some by up to ten seconds. (These delays may also arise from delays at the workstation. Lisp Machines are not noted for real-time response.)

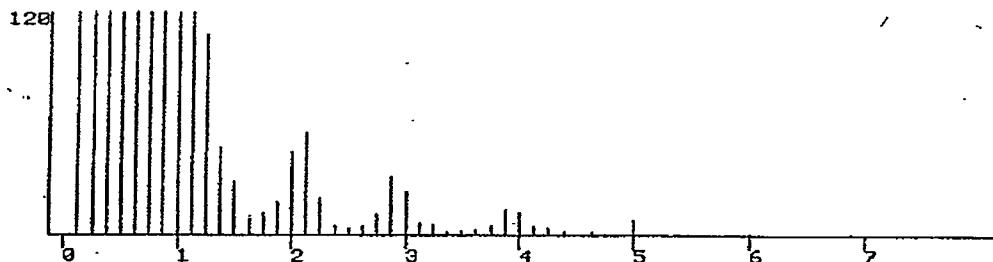


Figure C-4: Histogram of inter-arrival times of packets

It would be better to have a protocol which guarantees to deliver data intact and free of errors, if it delivers it at all, but does not guarantee to deliver the data. Real time data is only valuable in real time, and time spent retransmitting old data is taken away from never, more valuable data. Such a protocol modification is feasible for the Spectruin product (personal communication), and might be required for future work.

Index

access ramp 50
 ATM 88
 bear 27
 Bridge, the 151
 CARIN 107
 cellular phone 150
 confidence 67, 80
 dead reckoning 137
 Dectalk 101
 deictic 33
 DIME 42
 DriverGuide 105
 Elliot, R. J. 99
 ERGS 98
 ESA 135
 ETAK 56, 106
 EVA 57
 exit 67
 gas stations 88
 GDOP 130
 GEOSTAR 135
 GLONASS 134
 GPS 132
 graphs 99
 Grice, H. P. 35
 Hertz 105
 ice cream 88
 Kuipers, Benjamin 29
 lane advice 33
 legal connectivity 44
 Lesk, Michael 99
 LORAN 129
 Lynch, Kevin 32
 map matching 139
 modem 150
 NAVSAT 135
 NAVSTAR 132
 repeating 89
 rotary 49, 68
 signs 53
 spelling 102
 STARFIX 135
 street quality 49

TIGER 47

toll booths 54

traffic lights 53

TRANSIT 133

turn 25

turning radius 77

turning speed 77

U Turn 104

user model 66

warnings 90

wobble 45

Yellow Pages 88

Bibliography

- [1] Federal Highway Administration. Economic assessment of potential solutions for improving motorist route following. Technical Report RD-86/029, Federal Highway Administration, June 1986.
- [2] IEEE Aerospace and Electronic Systems Society. *IEEE Position and Location Symposium*. IEEE, 1986. IEEE reference 86CH2365-5.
- [3] IEEE Aerospace and Electronic Systems Society. *IEEE Position and Location Symposium*. IEEE, 1988. IEEE reference 88CH26757.
- [4] Ron Alexander. A Way to Go From A to B. *New York Times*, 1989. March 11, page 52.
- [5] Douglas E. Appelt. Some pragmatic issues in the planning of definite and indefinite noun phrases. In *Proceedings of the 23rd conference of the Association for Computational Linguistics*, pages 198–203, 1985.
- [6] G. A. Arrondo, J. C. Feggeler, and J. I. Smith. Voice and data transmission. *Bell System Technical Journal*, 58(1):97–122, Jan 1979.
- [7] Javad Ashjaee. Differential GPS with Ashtech XII. In *IEEE Position and Location Symposium*, pages 318–322, 1988. IEEE CH2675-7.
- [8] Herman J. G. M. Benning. Digital maps on compact discs. Technical Paper Series 860125, Society of Automotive Engineers, 1986.

- [9] Earl G. Blackwell. Overview of Differential GPS Methods. *Navigation: Journal of The Institute of Navigation*, 32(2):114–125, Summer 1985.
- [10] Matthew M. Blizzard. Update: Differential loran-c system. Technical report, USCG R&D Center, Groton, CT 06340, 1988.
- [11] A. Borelli and S. Sklar. An integrated approach to Automatic Vehicle Monitoring and Mobile Digital Communications. In *Proceedings of the 28th IEEE Vehicular Technology Conference*, March 1978. IEEE publication 78CH12971.
- [12] Mark Bosworth and Rodman Low. Taxi!: Dynamic cartographic software for training cab drivers. Technical report, Hunter College Department of Geology and Geography, (212)-772-4000, 1988. paper presented at the Annual Meeting of the Association of American Geographers.
- [13] Gerard Boudriault. Topology in the TIGER file. In *AUTO CARTO 8 International Symposium on Automation in Cartography*, April 1987.
- [14] Peter Braegas. Function, Equipment, and Field Testing of a Route Guidance and Information System for Drivers (ALI). *IEEE Transactions on Vehicular Technology*, VT-29(2):216–225, May 1980.
- [15] Donald F. Cooke. Vehicle Navigation Appliances. In *AUTO CARTO 7 International Symposium on Automation in Cartography*, pages 108–115, March 1985.
- [16] S. A. Dale and P. Daly. Recent Observations on the Soviet Union's Glonass Navigation Satellites. In *IEEE Position and Location Symposium*, pages 20–25, 1986.
- [17] S. A. Dale, I D. Kitching, and P. Daly. Position Fixing Using the USSR's Glonass C/A Code. In *IEEE Position and Location Symposium*, pages 13–20, 1988.

- [18] James R. Davis. A voice interface to a direction giving program. Technical Report 2, MIT Media Laboratory Speech Group, Apr 1988. Replaces the paper in the 1986 AVIOS proceedings titled "Giving Directions: A voice interface to an urban navigation program".
- [19] James R. Davis and Thomas F. Trobaugh. Direction assistance. Technical Report 1, MIT Media Laboratory Speech Group, Dec 1987.
- [20] Arthur R. Dennis. STARFIX. In *IEEE Position and Location Symposium*, pages 251–255, 1986.
- [21] R. J. Elliot and M. E. Lesk. Let your fingers do the driving: Maps, yellow pages, and shortest path algorithms. Technical Report unpublished, Bell Laboratories, 1982.
- [22] R. J. Elliot and M. E. Lesk. Route finding in street maps by computers and people. In *Proceedings of the National Conference on Artificial Intelligence*, pages 258–261, 1982.
- [23] Amr A. El-Sawy et al. LORAN C Tracking of Land Vehicles using micro computers. In *Proceedings of the 28th IEEE Vehicular Technology Conference*, pages 157–161, March 1978. IEEE publication 78CH12971.
- [24] Gregory M. Fontana. Transmission of Data over the Cellular Telephone Network. In *Proceedings of the 37th IEEE Vehicular Technology Conference*, pages 528–531, June 1987. IEEE publication 87CH24299.
- [25] Seiichi Fujimura et al. A vibration gyros and their applications. In *Digest of Technical Papers*, pages 116–117. IEEE ICCE, 1989. IEEE publication CH2724-3.
- [26] F. K. Ganjon. Liability aspects of electronic charts. In *IEEE Position and Location Symposium*, pages 303–305, 1986.

- [27] David K. Gifford, John M. Lucassen, and Stephen T. Berlin. The application of digital broadcast communication to large scale information systems. *IEEE Journal on Selected Areas in Communication*, 3(3):457–467, May 1985.
- [28] H. P. Grice. Logic and conversation. In Cole and Morgan, editors, *Syntax and Semantics: Speech Acts*, volume 3, pages 41–58. Academic Press, 1975.
- [29] George W. Gruver and Otto A. Reichardt. The Huntington Beach Automatic Vehicle Monitoring System using overlapping RF signposts. In *Proceedings of the 28th IEEE Vehicular Technology Conference*, pages 538–542, March 1978. IEEE publication 78CH12971.
- [30] Peter Haeussermann. On board computer system for navigation, orientation, and route optimization. Technical Paper Series 840483, Society of Automotive Engineers, 1984.
- [31] Robert G. Hajovsky. An Update on the Huntington Beach Automatic Vehicle Location System. In *Proceedings of the 31rst IEEE Vehicular Technology Conference*, April 1981. IEEE publication 81CH1638-6.
- [32] Herbert H. Hall. Alternate approaches to automobile navigation. Technical Paper Series 861057, Society of Automotive Engineers, 1986.
- [33] Clyde B. Harris, Laurie A. Klesh an Edward J. Krawkiwsky, Hassan A. Karimi, and Ness S. T. Lee. Digital Map Dependent Functions of Automatic Vehicle Location Systems. In *IEEE Position and Location Symposium*, pages 79–87, 1988. IEEE CH2675-7.
- [34] Ted Harris and David Kerridge. The geomagnetic storm of 13 March 1989. *Radio Communication*, 65(5), May 1989.
- [35] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on SSC*, 4:100–107, 1968.

- [36] Stanley K. Honey, Marvin S. White Jr., and Walter B. Zavoli. Extending Low Cost Land Navigation Into Systems Information Distribution and Control. In *IEEE Position and Location Symposium*, pages 439–444, 1986. IEEE 86CH2365-5.
- [37] Tom Hunter and Javad Ashjaee. Land Navigation and Fleet Management with GPS, LORAN, and dead reckoning sensors. In *IEEE Position and Location Symposium*, pages 54–60, 1988. IEEE CH2675-7.
- [38] Toshiyuki Itoh, Yasuhiko Okada, Akira Endoh, and Kenji Suzuki. Navigation systems using gps for vehicles. Technical Paper Series 861360, Society of Automotive Engineers, 1986.
- [39] Robert V. Janc. Consideration of the Various Error Sources in a Practical Automatic Vehicle Location system. In *Proceedings of the 34th IEEE Vehicular Technology Conference*, pages 277–284, 1984. IEEE publication number CH19513.
- [40] Rudolph M. Kalafus, Janis Vilcans, and Norman Knable. Differential Operation of NAVSTAR GPS. *Navigation: Journal of The Institute of Navigation*, 30(3):187–204, Fall 1983.
- [41] Gerald S. Kaplan and Andrew D. Ritzie. An X-Band System using Semi-passive Signpost Reflectors for Automatic Location and Tracking of Vehicles. *IEEE Transactions on Vehicular Technology*, VT-26(1):18–22, February 1977.
- [42] Christine Kinnear. The TIGER Structure. In *AUTO CARTO 8 International Symposium on Automation in Cartography*, April 1987.
- [43] M. D. Kotzin and A. P. van den Heuvel. Dead Reckoning Vehicle Location using a solid state rate gyro. In *Proceedings of the 31rst IEEE Vehicular Technology Conference*, pages 169–172, April 1981. IEEE publication 81CH1638-6.

- [44] E. J. Krakiwsky, C. B. Harris, and R. V. C. Wong. A Kalman Filter for Integrating Dead Reckoning, Map Matching, and GPS Positioning. In *IEEE Position and Location Symposium*, pages 39–46, 1988. IEEE CH2675-7.
- [45] L. R. Kruczynski et al. Global positioning system differential navigation tests at the yuma proving ground. *Navigation: Journal of The Institute of Navigation*, 32(2):126–138, Summer 1985.
- [46] Benjamin J. Kuipers. *Representing Knowledge of Large-Scale Space*. PhD thesis, MIT, July 1977. Issued as Technical Report 418.
- [47] Benjamin J. Kuipers. Modelling human knowledge of routes: Partial knowledge and individual variation. In *Proceedings of the National Conference on Artificial Intelligence*, pages 216–219, 1983.
- [48] W. Thompson Lawrence. A Magnetic Signpost AVM System with Limited Dead Reckoning. *IEEE Transactions on Vehicular Technology*, VT-26(1):23–29, February 1977.
- [49] M. Liberman and J. Pierrehumbert. Intonational invariance under changes in pitch range and length. In *Language Sound Structure*, chapter 10. MIT Press, 1981.
- [50] Kevin Lynch. *The Image of the City*. MIT Press, 1960.
- [51] Peeder Ma. An algorithm to generate verbal instructions for vehicle navigation using a geographic database. *The East Lakes Geographer*, 22:44–60, 1987.
- [52] Robert W. Marx. The TIGER system: automating the Geographic Structure of the United States Census. *Government Publications Review*, 13:181–201, 1986.
- [53] Robert W. Marx. GIS, TIGER, and Other Useful Acronyms. In *National Conference of Geographic Information Systems*. Canadian Institute for Surveying and Mapping, March 1989.

- [54] Robert W. Marx and Alan J. Saalfeld. Programs for assuring map quality at the bureau of the census. In *Fourth Annual Research Conference, Geography Division, Room 3203-4, Bureau of the Census, U. S. Department of Commerce, Washington DC 20233*, March 1988.
- [55] Clare D. McGillem et al. Experimentally determined accuracy and stability of loran c signals for land vehicle location. *IEEE Transactions on Vehicular Technology*, VT-31(1):15-21, Feb 1982.
- [56] Matthew McGranaghan, David M. Mark, and Michael D. Gould. Automated provision of navigation assistance to drivers. *The American Cartographer*, 14(2):121-138, 1987.
- [57] Douglas L. Milliken. personal communication, Dec 1988.
- [58] Ernst Peter Neukirchner and Wolf Zechnall. Digital Map Data Bases for Autonomous Vehicle Navigation Systems. In *IEEE Position and Location Symposium*, pages 320-324, 1986. IEEE 86CH2365-5.
- [59] Osamu Ono, Hidemi Ooe, and Masahiro Sakamoto. CD-ROM Assisted Navigation System. In *Digest of Technical Papers*, pages 118-119. IEEE ICCE, 1988.
- [60] Harry M. O'Sullivan. Cellular telephone data communication system and method. United States Patent 4,697,281, September 1987.
- [61] Lee Ott. STARFIX: Commercial satellite positioning. In *IEEE Position and Location Symposium*, pages 8-12, 1988.
- [62] Georges Pachiaudi, Monique Vernet, and Anne Foti. Noise and speech interference in light automotive vehicles. Technical Paper Series 885118, Society of Automotive Engineers, 1988..
- [63] Cécile L. Paris. Tailoring object descriptions to a user's level of expertise. *Computational Linguistics*, 14(3):64-78, September 1988.

- [64] Timothy J. Peters. Automobile navigation using a magnetic flux gate compass. *IEEE Transactions on Vehicular Technology*, VT-35(2):41–47, May 1986.
- [65] Otmar Pilsak. Eva: An electronic traffic pilot for motorists. Technical Paper Series 860346, Society of Automotive Engineers, 1986.
- [66] Alexander H. Riccio. Automatic Vehicle Location Dallas Test Results. In *Proceedings of the 27th IEEE Vehicular Technology Conference*, pages 144–152, March 1977. IEEE publication 77CH11767.
- [67] Robert T. Richards and Leslie O. Snively. GEOSTAR positioning analysis. In *IEEE Position and Location Symposium*, pages 13–19, 1986.
- [68] Christopher K. Riesbeck. “You Can’t Miss It”: Judging the Clarity of Directions. *Cognitive Science*, 4:285–303, 1980.
- [69] Mortimer Rogoff. Electronic Charts - At the Heart of 21st Century Navigation. In *IEEE Position and Location Symposium*, pages 88–94, 1988. IEEE CH2675-7.
- [70] Dan A. Rosen, Frank J. Mammano, and Rinaldo Favout. An electronic route-guidance system for highway vehicles. *IEEE Transactions on Vehicular Technology*, VT-19(1):143–152, Feb 1970.
- [71] Ronald Rosenberg. Mapping out a new idea. *The Boston Globe*, 1987. February 17, page 39.
- [72] Harvey Sacks, Emanuel A. Schegloff, and Gail Jefferson. A simplest systematics for the organization of turn-taking for conversation. *Language*, 50(4):696–735, 1974. Reprinted in *Studies in the Organization of Conversational Interaction*, J. Schenken, ed., Academic Press 1978.
- [73] Richard C. Sagers. A LORAN-C based receiver for automatic vehicle location. In *Proceedings of the 35th IEEE Vehicular Technology Conference*, pages 387–391, 1986. IEEE CH2308-5.

- [74] Stephen R. Sampson. A survey of commercially available positioning systems. *Navigation: Journal of The Institute of Navigation*, 32(2):139–148, Summer 1985.
- [75] C. Schmandt and B. Arons. A conversational telephone messaging system. *IEEE Trans. on Consumer Electr.*, CE-30(3):xxi–xxiv, 1984.
- [76] Chris Schmandt. Employing voice back channels to facilitate audio document retrieval. In *Proceedings*, pages 213–218. ACM Conference on Office Information Systems, 1988.
- [77] Christopher Schmandt. Speech synthesis gives voiced access to an electronic mail system. *Speech Technology*, 2(3):66–69, 1984.
- [78] A. J. Sedlock. Mid-continent LORAN-C expansion. In *IEEE Position and Location Symposium*, pages 520–524, 1986.
- [79] Joel Sobel. Principal components of the census bureau's tiger file. Research in Contemporary and Applied Geography Discussion Series 3, Department of Geography, SUNY Binghamton, 1986.
- [80] G. J. Sonnenberg. *Radar and Electronic Navigation*. Butterworths, sixth edition, 1988.
- [81] Lynn A. Streeter and Diane Vitello. A profile of drivers' map reading abilities. *Human Factors*, 28:223–239, 1986.
- [82] Lynn A. Streeter, Diane Vitello, and Susan A. Wonsiewicz. How to tell people where to go: comparing navigational aids. *International Journal of Man/Machine Systems*, 22(5):549–562, May 1985.
- [83] Katsutoshi Tagami et al. Electro gyro-cator: New inertial navigation system for use in automobiles. Technical Paper Series 830659, Society of Automotive Engineers, 1983.

- [84] C. Tarriere et.al. Some ergonomic features of the driver vehicle environment interface. Technical Paper Series 885051, Society of Automotive Engineers, 1988.
- [85] Martin L. G. Thoone and Ron M. A. M. Breukers. Applications of the compact disc in car information and navigation systems. Technical Paper Series 840156, Society of Automotive Engineers, 1984.
- [86] T. Tsumura, M. Hashimoto, and N. Fujiwara. A Vehicle Position and Heading Measurement System Using Corner Cube and Laser Beam. In *IEEE Position and Location Symposium*, pages 47–53, 1988.
- [87] Lawrence A. Whitcomb. Using Low Cost Magnetic Sensors on Magnetically Hostile Land Vehicles. In *IEEE Position and Location Symposium*, pages 54–60, 1988. IEEE CH2675-7.
- [88] Walter W. Wierwille et al. Strategic use of visual resources by the driver while navigating with an in-car navigation display system. Technical Paper Series 885180, Society of Automotive Engineers, 1988.
- [89] Walter B. Zavoli and Stanley K. Honey. Map Matching Augmented Dead Reckoning. In *Proceedings of the 35th IEEE Vehicular Technology Conference*, pages 359–362, 1986. IEEE CH2308-5.



UNITED STATES DEPARTMENT OF COMMERCE
Patent and Trademark Office

Address : COMMISSIONER OF PATENTS AND TRADEMARKS
 Washington, D.C. 20231

SERIAL NUMBER	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.
07/565,274	05/09/90	DAVIS	J

EXAMINER
PIPALA, E

SAM PASTERNACK
 CHOATE, HALL & STEWART
 EXCHANGE PLACE
 53 STATE STREET
 BOSTON, MA 02109

ART UNIT	PAPER NUMBER
2304	5

DATE MAILED: 11/08/91

This is a communication from the examiner in charge of your application.
 COMMISSIONER OF PATENTS AND TRADEMARKS

This application has been examined Responsive to communication filed on _____ This action is made final.

A shortened statutory period for response to this action is set to expire 3 month(s), — days from the date of this letter.
 Failure to respond within the period for response will cause the application to become abandoned. 35 U.S.C. 133

Part I THE FOLLOWING ATTACHMENT(S) ARE PART OF THIS ACTION:

1. Notice of References Cited by Examiner, PTO-892 (2 - sheets)
2. Notice re Patent Drawing, PTO-948.
3. Notice of Art Cited by Applicant, PTO-1449. (14 - sheets)
4. Notice of Informal Patent Application, Form PTO-152
5. Information on How to Effect Drawing Changes, PTO-1474.
6. _____

Part II SUMMARY OF ACTION

1. Claims 1 through 58 are pending in the application.

Of the above, claims _____ are withdrawn from consideration.

2. Claims _____ have been cancelled.

3. Claims _____ are allowed.

4. Claims 1 through 58 are rejected.

5. Claims _____ are objected to.

6. Claims _____ are subject to restriction or election requirement.

7. This application has been filed with informal drawings under 37 C.F.R. 1.85 which are acceptable for examination purposes.

8. Formal drawings are required in response to this Office action.

9. The corrected or substitute drawings have been received on _____. Under 37 C.F.R. 1.84 these drawings are acceptable; not acceptable (see explanation or Notice re Patent Drawing, PTO-948).

10. The proposed additional or substitute sheet(s) of drawings, filed on _____, has (have) been approved by the examiner; disapproved by the examiner (see explanation).

11. The proposed drawing correction, filed _____, has been approved; disapproved (see explanation).

12. Acknowledgement is made of the claim for priority under U.S.C. 119. The certified copy has been received not been received been filed in parent application, serial no. _____; filed on _____.

13. Since this application appears to be in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under Ex parte Quayle, 1935 C.D. 11; 453 O.G. 213.

14. Other

0438

EXAMINER'S ACTION

Art Unit 234

Ser. No. 07/565,274

1

The title of the invention is not descriptive. A new title is required that is clearly indicative of the invention to which the claims are directed.

The drawings are objected to under 37 CFR 1.83(a). The drawings must show every feature of the invention specified in the claims. Therefore, the line segments and nodes of claim 2, the use of DIME and TIGER files and the added extensions of claims 3-6, the representation of a three-dimensional topology, use of street quality, divided streets, signs ,traffic lights, landmarks and buildings, lane information, speed limits, expected rate of travel, et cetera, of claims 7-20 and 23, and the updating of the database by radio broadcast of claims 21 and 22, a flowchart or list of the different algorithm(s), routine(s) and parameters used as part of the route finder and the new route capability of claims 24-31, some representation of the dead-reckoning means as far as the use of an odometer, magnetic compass, steering wheel sensor, differential odometer, gyroscope as well as map matching and the use of a hybrid system (claims 32-40), the different aspects of the discourse generator and speech functions as found in claims 41-47 and 51-56, the different aspects of the driver input means found in claims 48-50, and finally the use of a remote computing apparatus and the use of two cellular phones for transmission of a data channel and a voice channel. These features must be shown or the feature

Art Unit 234

Ser. No. 07/565,274

2

should be cancelled from the claim(s). No new matter should be entered. Even though the above listed deficiencies in the drawings relate to the operations of the route finder, location system and discourse generator and other subsystems shown in Fig. 1, the above are nevertheless readily adapted to description by means of a flow chart or by means of a block diagram as directed at the end of 1.83(a).

The disclosure is objected to because of the following informalities: p. 11, l. 5, "streets" should be "streets"; p. 25, l. 17, change "it" to "is"; p. 25, l. 30, "hearing" is misspelled; on p. 34, l. 16, the phrase "pertain the driving" seems misplaced; p. 35, l. 9, please replace "beable" with "able"; p. 38, l. 15 (or so), please change "The minimum time that is can estimate" to "The minimum time that it can estimate"; p. 39, l. 6, please replace "is" with "in"; p. 40, l. 16, please change "Give that the driver" to "Given that the driver"; p. 44, l. 2, "never" should be "ever". Appropriate correction is required.

With respect to applicant's Information Disclosure Statement, only those references of which copies have been received have been initialed by the examiner as having been considered. This is according to common practice as specified by 37 CFR 1.98 and MPEP 609. The patent to O'Sullivan was loc~~0440~~

Art Unit 234

Ser. No. 07/565,274

3

by the examiner and is listed on the Notice of References Cited (PTO-892), so that applicant need not order this particular piece of art. However, many of the prior art documents cited by applicant are seen to be extremely relevant to the invention given their description in the specification, in Mr. Davis' Ph.D. thesis, and in the Information Disclosure Statement. These documents may be used as basis for all or part of a rejection, as needed, and applicant is requested to forward any such prior art documents (those used in a rejection) along with any others that applicant would like to have considered by the examiner and thereby formally made of record. Attached to or enclosed with this Office Action applicant will find copies of all fourteen (14) pages of PTO-1449 listing the documents which were cited in the disclosure, the examiner has initialed and considered those documents for which copies were enclosed. Those documents for which there was no copy were crossed from the list(s) with a line. For the sake of limiting the number of sheets of PTO-1449 the examiner has only used a pencil on the Office's copies, so that if copies of the documents are forwarded then it will not be necessary to re-list them on additional copies of PTO-1449.

Claims 1-58 are rejected under 35 U.S.C. § 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

0441

Art Unit 234

Ser. No. 07/565,274

4

The claims are seen to be vague and indefinite in that, apart from a limited preamble and the recitation of the discourse generator, there is nothing connecting the different elements of the invention with each other in a manner which defines a working automobile navigation system. As it is claim 1 seems to merely recite a number of elements (computing apparatus, driver input means, map database, position sensing, location sensing, route finding, discourse generation and speech generation), without enabling one to understand the manner in which these elements are to interact with one another. In this manner the recitation of these elements is aggregative in nature without a clearly recited connection/interaction between them.

The following is a quotation of the appropriate paragraphs of 35 U.S.C. § 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless --

(e) the invention was described in a patent granted on an application for patent by another filed in the United States before the invention thereof by the applicant for patent, or on an international application by another who has fulfilled the requirements of paragraphs (1), (2), and (4) of section 371(c) of this title before the invention thereof by the applicant for patent.

Claims 1-58 are rejected under 35 U.S.C. § 102(e) as being clearly anticipated by the Ph.D. thesis of James Raymond Davis.

Art Unit 234

Ser. No. 07/565,274

5

The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

Nakahara et al. (P.N. 5,041,983) teaches a guidance system for a vehicle which determines the shortest route between a starting point and an ending point entered by a driver, and gives spoken directions for a driver to follow.

Champion, III et al. (P.N. 4,812,843) teaches a traffic information system in which individual subscribers use telephone, mobile telephone or computer to receive updated reports of traffic conditions, specific routings, and/or other information which may be of interest to the subscriber. In col. 9, ll. 42-46, the use of a telephone keypad is taught for entering subscriber coordinates.

Yamada (P.N. 4,926,336) teaches a route searching system in which weighted information on road data is used to distinguish between possible routes. Figures 2C and 4 both teach aspects of physical and legal connectivity.

Neukrichner et al. (P.N. 4,984,168) teach some of the aspects of the Bosch "EVA" system which provides route searching and trip guidance to a driver. This reference also teaches factoring in road quality as resistance values, dead-reckoning equipment and route changing/error options (col. 6).

Nimura et al. (P.N. 4,992,947) disclose a vehicle navigation system where the main feature is the "HELP" function, but which

0443

Art Unit 234

Ser. No. 07/565,274

6

also teaches the use of physical and legal connectivity, stoplights, landmarks, segments, nodes and their coordinates, named intersections and other types of information for a route.

Yokoyama et al. (P.N. 5,043,902) teach another vehicle navigation system in which physical and legal connectivity is considered in the route search, where vehicle position is determined by means of distance and steering angle sensors, and in which the guidance information is given by a voice output unit.

Savage et al. (P.N. 4,954,958) teach a direction information system in which a central computing station provides routing and travel directions between two locations through any of a number of different communications means. Column 2, lines 4-14, teach the availability of a number of different route types or parameters, where column 4 teaches the use of digitized or synthesized voice in conjunction with an English language direction format, as well as U.S. Geological Service digitized maps and other data from an on-line directory listing database. Columns 11 and 12 disclose the type of programming language used, the use of a portable transceiver such as a mobile phone and modem, and the feature of a lost-unlost function.

Gray et al. (P.N. 4,891,761) teach a method for updating a digitized map, such as the DIME type developed by the Bureau of Census, by means of a LORAN-C position determining system and a second channel for voice contact with the central station.

0444

Art Unit 234

Ser. No. 07/565,274

7

Ross et al. (P.N. 5,021,961) teach a data processing device which provides data relating to categories of services available along a route, such as restaurants, gas stations, motels and the like.

Nimura et al. (P.N. 4,939,662) disclose a navigation system for a vehicle in which the distance between instructions given to a driver is such that the driver is unfailingly supplied with information with respect to turns, told about distinguishing features along the course, and frequently supplied with guidance information so as to assure the driver or the vehicle.

Honey et al. (P.N. 4,796,191) seems to be one of the patents covering the "ETAK" type vehicle navigation system, wherein the detected distance and heading of a vehicle is determined providing a dead reckoned position which is compared with an data from a stored map database so as to determine the most probable position of the vehicle. In col. 6, ll. 47-56, the use of a flux gate compass and differential wheel sensors (a differential odometer) is taught, where col. 8 teaches some of the database map enhancements such as street widths and vertical slopes of line segments (making for a 3-D representation of the map).

Honey et al. (P.N. 4,734,863) is another reference which teaches the map matching used by the "ETAK" system.

De Villeroche (P.N. 4,951,211) teaches an electronic guidance and information system which is capable of route selection, physical and legal connectivity (col. 1, ll. 32-44) 0445

Art Unit 234

Ser. No. 07/565,274

8

stores the names of roads, sites and nearby restaurants, and is capable of selecting a new route from a new starting point. In col. 6, ll. 18-41, additional information such as driving customs, traffic signals, road safety, sight seeing, special information and warnings are given to the user.

Ingels (P.N. 4,139,889) teaches a vehicle position indication system which transmits distance and direction information in a digital format to a central location.

Nimura et al. (4,937,751) disclose a vehicle navigation apparatus which provides spoken guidance information for a number of different characterizing features along the road, such as bridges, railroad crossings, underpasses, tunnels, police stations, post offices and the like.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Edward Pipala whose telephone number is (703) 308-1397.

Any inquiry of a general nature or relating to the status of this application should be directed to the Group receptionist whose telephone number is (703) 308-0754.

EP
Edward Pipala

November 4, 1991

(703) 308-1397

Mal
PARBUDAM S. LALL
SUPERVISORY PATENT EXAMINER
ART UNIT 234

0446

PTO - 948
(Rev. 8-62)U.S. DEPARTMENT OF COMMERCE
PATENT AND TRADEMARK OFFICEATTACHMENT TO
PAPER NUMBER

5

S.N.

563274

GROUP 234

PTO draftsmen now review both

formal and informal drawings.

The examiner will require

submission of formal drawings at the appropriate time.

Drawing Corrections and/or new drawings may only be submitted in the manner set forth in the attached letter, "Information on How to Effect Drawing Changes" PTO-1474.

A. The drawings, filed on 8/9/20, are objected to as informal for reason(s) checked below:

1. <input type="checkbox"/> Lines Pale.	11. <input type="checkbox"/> Parts in Section Must Be Hatched.
2. <input type="checkbox"/> Paper Poor.	12. <input type="checkbox"/> Solid Black Objectionable.
3. <input checked="" type="checkbox"/> Numerals Poor / small must be at least <i>fig 1,2</i> <i>1/8" tall</i>	13. <input checked="" type="checkbox"/> Figures Legends <i>Poor</i> <i>poor</i> <i>poorly</i> .
4. <input checked="" type="checkbox"/> Lines Rough and Blurred. <i>fig 1,2</i>	14. <input type="checkbox"/> Mounted Photographs.
5. <input type="checkbox"/> Shade Lines Required.	15. <input type="checkbox"/> Extraneous Matter Objectionable. [37 CFR 1.84 (1)]
6. <input type="checkbox"/> Figures Must be Numbered.	16. <input checked="" type="checkbox"/> Paper Undersized; either 8½" x 14", or 21.0 cm. x 29.7 cm. required.
7. <input type="checkbox"/> Heading Space Required.	17. <input type="checkbox"/> Proprietary Margin Required: <input type="checkbox"/> TOP 2.5 cm. <input type="checkbox"/> RIGHT 1.5 cm. <input type="checkbox"/> LEFT 2.5 cm. <input type="checkbox"/> BOTTOM 1.0 cm.
8. <input type="checkbox"/> Figures Must Not be Connected.	
9. <input type="checkbox"/> Criss-Cross Hatching Objectionable.	
10. <input type="checkbox"/> Double-Line Hatching Objectionable.	18. <input type="checkbox"/> Other:

B. The drawings, submitted on 8/9/20, are so informal they cannot be corrected. New drawings are required. Submission of the new drawings MUST be 0447 made in accordance with the attached letter.

TO SEPARATE, HOLD TOP AND BOTTOM EDGES, SNAP-APART AND DISCARD CARBON

(A)

1 of 2				SERIAL NO.	GROUP/ART UNIT	ATTACHMENT TO PAPER NUMBER		
				07/565,274	234	5		
NOTICE OF REFERENCES CITED				APPLICANT(S)	James R. Davis et al. (1/2)			
U.S. PATENT DOCUMENTS								
*	DOCUMENT NO.	DATE	NAME	CLASS	SUB-CLASS	FILING DATE IF APPROPRIATE		
A	4697281	9/87	O'Sullivan	455	33(x)			
B	5041983	8/91	Nakahara et al.	364	449	3/30/90		
C	4812843	3/89	Champion, III et al.	340	989(x)			
D	4926336	5/90	Yamada	364	444			
E	4984168	1/91	Neukrichner et al.	364	449	8/24/89		
F	4992947	2/91	Nimura et al.	364	444	2/27/88		
G	5043902	8/91	Yokoyama et al.	364	449	8/23/89		
H	4954958	9/90	Savage et al.	364	444	8/19/88		
I	4891761	1/90	Gray et al.	364	449(x)	3/31/88		
J	5021961	6/91	Ross et al.	340	990(x)	9/21/89		
K	4939662	7/90	Nimura et al.	340	990(x)			
FOREIGN PATENT DOCUMENTS								
*	DOCUMENT NO.	DATE	COUNTRY	NAME	CLASS	SUB-CLASS	PERTINENT SHTS. DWG	PP. SPEC.
L								
M								
N								
O								
P								
Q								
OTHER REFERENCES (Including Author, Title, Date, Pertinent Pages, Etc.)								
R								
S								
T								
U								
EXAMINER	Edward Pipala	DATE	11/4/91					
* A copy of this reference is not being furnished with this office action. (See Manual of Patent Examining Procedure, section 707.05 (a).)								

0448

EXHIBIT

2

(Part 9)

TO SEPARATE, HOLD TOP AND BOTTOM EDGES, SNAP-APART AND DISCARD CARBON

CJY

FORM PTO-892 (REV. 3-78)	U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE	SERIAL NO. 07/565,274	GROUP ART UNIT 234	ATTACHMENT TO PAPER NUMBER 5
NOTICE OF REFERENCES CITED		APPLICANT(S) James R. Davis et al.		

U.S. PATENT DOCUMENTS

*	DOCUMENT NO.	DATE	NAME	CLASS	SUB-CLASS	FILING DATE IF APPROPRIATE
A	4796191	1/89	Honey et al.	364	450	
B	4951211	8/90	De Villerche	364	444	4/28/89
C	4139889	2/79	Ingels	340	989.65	
D	4937751	6/90	Nimura et al.	364	444(6)	
E	4734863	3/88	Honey et al.	340	988(6)	
F						
G						
H						
I						
J						
K						

FOREIGN PATENT DOCUMENTS

*	DOCUMENT NO.	DATE	COUNTRY	NAME	CLASS	SUB-CLASS	PERTINENT SHTS. PP. DWG. SPEC.
L							
M							
N							
O							
P							
Q							

OTHER REFERENCES (Including Author, Title, Date, Pertinent Pages, Etc.)

R	
S	
T	
U	

EXAMINER <i>Edward Pipala</i>	DATE <i>11/4/91</i>	
* A copy of this reference is not being furnished with this office action. (See Manual of Patent Examining Procedure, section 707.05 (a).)		

0636

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: J.R. Davis and C.M. Schmandt Examiner: E. Pipala
Serial No. 07565,274 Art Unit: 2304
Filed: August 9, 1990
For: AUTOMOBILE NAVIGATION SYSTEM



Commissioner of Patents and Trademarks
Washington, D.C. 20231

H
G
RECEIVED

MAY 14 1991

Sir:

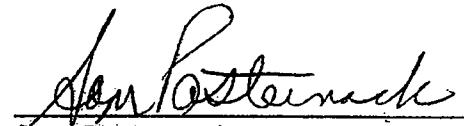
GROUP 230

PETITION FOR EXTENSION OF TIME

A three (3) month extension of time is requested to respond to the Office Action mailed November 8, 1991. This extension will extend the period of response through May 8, 1992. Accordingly, a check for \$405.00 is enclosed in view of the small entity status of the application.

Please charge any additional fees which may be required to Deposit Account No. 03-1721.

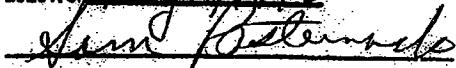
Respectfully submitted,


Sam Pasternack
Reg. No. 29,576

Choate, Hall & Stewart
Exchange Place
53 State Street
Boston, MA 02109

May 5, 1991

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner of Patents and Trademarks, Washington, D.C. 20231 on May 5, 1991.



120 BA 05/13/92 07565274

1 217 405.00 CK

0798



71A

RECEIVED

MAY 14 1991

GROUP 230

6-4

W
6/14/92
Applicant: J.R. Davis and C.M. Schmandt
Serial No.: 565,274
Filed: August 9, 1990
For: Automobile Navigation System

Art Unit: 2304
Examiner: E. Pipala

Commissioner of Patents and Trademarks
Washington, D.C. 20231

Amendment

In response to the Office Action mailed November 8, 1991, please amend the application as follows:

In the Title:

Replace "Automobile Navigation System" with - - Automobile Navigation
System Using Real Time Spoken Driving Instructions - -

In the Specification:

Page 4, after line 12, insert:

- - Fig. 3 is a schematic illustration of the map database.

Fig. 4 is a schematic illustration of the route finder.

Fig. 5 is a schematic illustration of the position sensor. - -

Page 11, line 5: replace "street" with - - streets - -

Page 11, after line 19, insert the following new paragraph:

- - The map database is shown schematically in Fig. 3. In the preferred embodiment, the map database 14 includes, as its basis, a file 28 of segments and

A3
nodes. File 28 may be an original file or may be adapted from a DIME file or a TIGER file by adding the above-described extensions. In addition, the map database 14 may include optional features 30, as described above. --

A4
Page 16, after line 2, insert the following new paragraph:

A4
The route finder is shown schematically in Fig. 4. In the preferred embodiment, the route finder 16 includes, as its basis, an algorithm 32. Algorithm 32 may be, for example, an original algorithm based on a best-first search algorithm, the A* algorithm, or a modified A* algorithm. In preferred embodiments, the route finder is adapted to find the best route according to any one of three cost metrics 34: distance, speed, simplicity. The route finder calculates a new route in the event of driver error or unforeseen circumstances, as indicated. --

A5
Page 17, after line 27, insert the following new paragraph:

A5
The position sensor is shown schematically in Fig. 5. As indicated, it includes a displacement sensor 36 and a direction sensor 38. --

A6
Page 25, line 17: replace "it" with -- is --.

Page 25, line 30: replace "healing" with -- hearing --.

Page 34, line 16: replace "the" with -- to --.

Page 35, line 7: replace "beable" with -- be able --.

Page 38, line 15: replace "is" with -- it --.

Page 39, line 6: replace "is" with -- in --.

Page 40, line 16: replace "Give" with -- Given --.

Page 44, line 2: replace "never" with -- ever --.

In the Claims:**Claim 1 (Amended)**

1. An automobile navigation system which [uses] produces spoken instructions to direct a driver of an automobile to a destination in real time comprising:

(1) computing apparatus [adapted to run and coordinate] for running and coordinating system processes,

(2) driver input means functionally connected to said computing apparatus [whereby the driver can enter] for entering data into said computing apparatus, said data including a desired destination,

(3) a map database functionally connected to said computing apparatus which distinguishes between physical and legal connectivity,

(4) position sensing apparatus installed in the automobile and functionally connected to said computing apparatus for providing said computing apparatus data for determining the automobile's current position,

(5) a location system functionally connected to said computing apparatus [which determines] for accepting data from said position sensing apparatus, for consulting said map database, and for determining the automobile's current position [on a map] relative to the map database [from data from the position sensing apparatus],

(6) a route-finder functionally connected to said computing apparatus, [which] for accepting the desired destination from said driver input means and the current position from said location system, for consulting said map database, and [computes] for computing a route to the destination [from any current position],

(7) a discourse generator functionally connected to said computing apparatus

46

[which composes] for accepting the current position from said location system and the route from said route finder, for consulting said map database, and for composing discourse including instructions and other messages [based on data from said location system, said route-finder, and said map database] for directing the driver to the destination from the current position,

A6
 P1 a speech generator functionally connected to said discourse generator [which generates] for generating speech from said discourse provided by said discourse generator, and

P1 voice apparatus functionally connected to said speech generator for communicating said speech provided by said speech generator to said driver.

In the Figures:

Please add the enclosed Figs. 3-5.

Remarks

Reexamination and reconsideration of the rejections are hereby requested for the following reasons.

The examiner has objected to the title for being not descriptive. In response, applicant has changed the title to "Automobile Navigation System Using Real Time Spoken Driving Instructions."

The examiner has objected to the drawings under 35 CFR 1.83(a) for not showing every feature of the invention specified in the claims. In a telephone interview, the examiner stated that simple block diagrams of the features amenable to illustration would suffice. In response, applicant has added Figs. 3-5, which illustrate the features claimed. No new matter has been added by this amendment.

The examiner has objected to the disclosure because of informalities. The typographical errors cited by the examiner have been corrected in the above amendments.

47

The examiner has rejected claims 1-58 under 35 U.S.C. 112, second paragraph, as being indefinite. Claim 1 has been amended to more particularly point out the connections and interactions between the different elements of the invention, as required by the examiner.

The examiner has rejected claims 1-58 under 35 U.S.C. 102(e) as being anticipated by the Ph.D. thesis of J.R. Davis. During a telephone conversation with the examiner, the examiner stated that the reason for the rejection was that the title page of the thesis bears a submission date of August 4, 1989, more than one year before the filing date of the present application. August 4 is the date that the thesis was signed, and not the date on which the thesis became available to the public. M.I.T. does not generally catalog and shelve theses until several months after the official date of submission. Enclosed is a copy of the title page of the M.I.T. library's copy of the thesis, which bears a date of February 27, 1990. Therefore, the thesis did not become available to the public more than a year before the filing date of the present application, and is therefore not 102 art with respect to the present application.

In response to the examiner's request, copies of references which were included in the Information Disclosure Statement filed with the application which the applicant considers pertinent to the present invention as claimed and which the applicant would like to be considered and made of record are enclosed and included on a new PTO-1449.

It is respectfully submitted that the claims are now in condition for allowance, and it is requested that a Notice of Allowance be issued.

Please charge any fees in connection with this response to our Deposit Account No. 03-1721.

Respectfully Submitted,

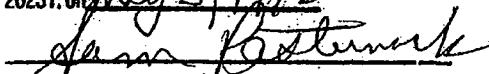


Sam Pasternack, Esq.
Reg. No. 29,576

Choate, Hall & Stewart
Exchange Place
53 State Street
Boston, MA 02109
(617) 227-5020

May 5, 1992

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner of Patents and Trademarks, Washington, D.C. 20231 on May 5, 1992



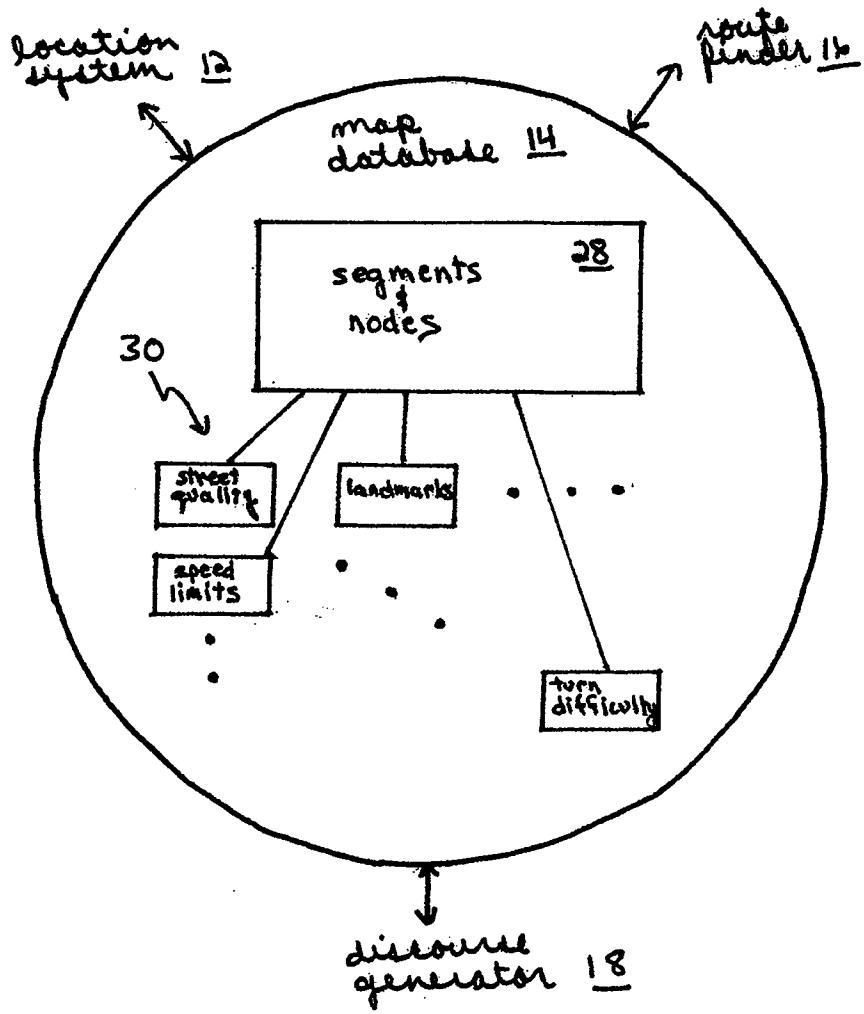


Fig. 3

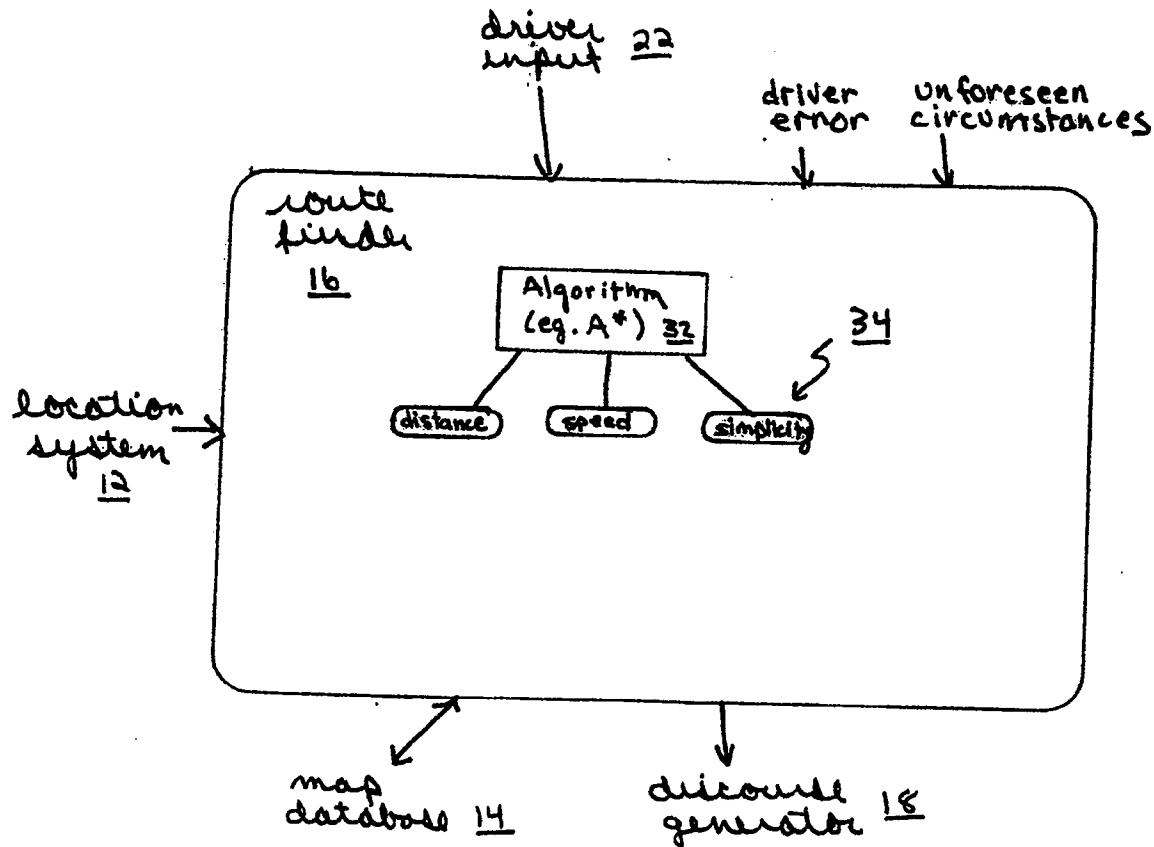


Fig. 4

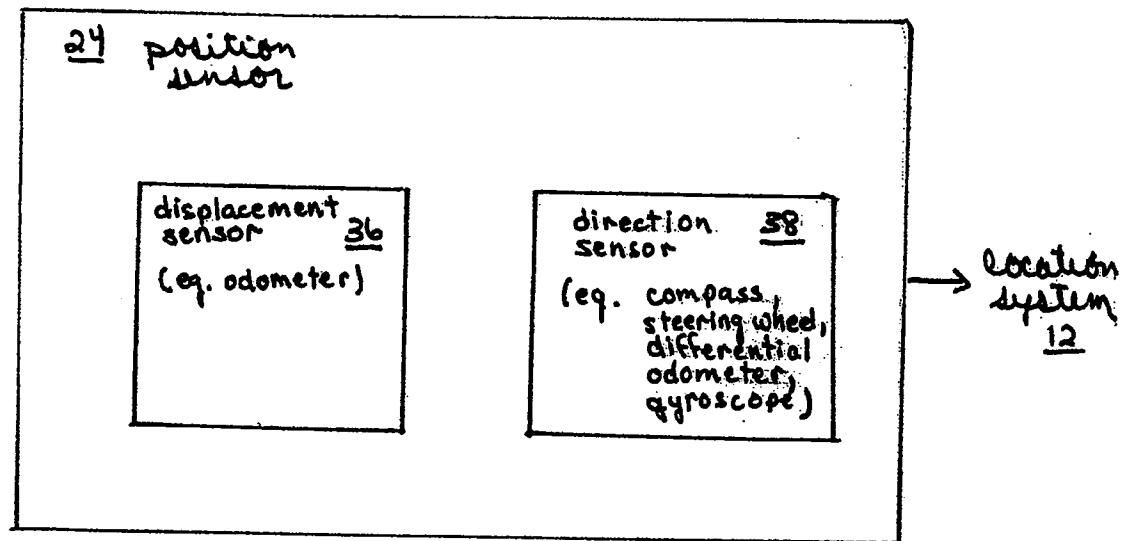


Fig. 5

Form PTO-1449
(REV. 8-84)U.S. Department of Commerce
Patent and Trademark OfficeAtty.
Docket
MITS088S 07/565,274**INFORMATION DISCLOSURE STATEMENT**

(Use several sheets if necessary)

Applicant
Davis et al.Filing
Date
8/9/90

Group

2304

U. S. PATENT DOCUMENTS

Examiner's initials	U.S. Patent No.	Applicant	Issue Date	C.P.	S.b.
lf	4,882,696	Nimura et al.	November 21 1989	-	-
gf	4,926,336	Yamada	May 15, 1990	-	-

FOREIGN PATENT DOCUMENTS

Examiner's initials	Document No.	Country	Date	Translation	
				Yes	No

Examiner's initials

OTHER DOCUMENTS
(Including Author, Title, Date, Pertinent Pages, Etc.)

EXAMINER

Edward Pipala

DATE CONSIDERED

6-23-92

EXAMINER: Initial if citation considered, whether or not citation is in conformance with MPEP 609; Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to applicant.



UNITED STATES DEPARTMENT OF COMMERCE
Patent and Trademark Office
Address: COMMISSIONER OF PATENTS AND TRADEMARKS
Washington, D.C. 20231

ATTORNEY/CLIENT NUMBER	FILED NUMBER	DATE RECEIVED	ATTORNEY DOCKET NO.
02/565,274	08/09/90	DAVIS	
		EXAMINER	
		PTIPALADE	
		ARTICLE	PAPER NUMBER
		2304	8
		DATE MAILED	06/30/92

NOTICE OF ALLOWABILITY

PART I.

- This communication is responsive to amendment filed 5-7-92.
- All the claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice Of Allowance And Issue Fee Due or other appropriate communication will be sent in due course.
- The allowed claims are 1-5B.
- The drawings filed on _____ are acceptable.
- Acknowledgment is made of the claim for priority under 35 U.S.C. 119. The certified copy has [...] been received. [...] not been received. [...] been filed in parent application Serial No. _____ filed on _____.
- Note the attached Examiner's Amendment.
- Note the attached Examiner Interview Summary Record, PTOL-413.
- Note the attached Examiner's Statement of Reasons for Allowance.
- Note the attached NOTICE OF REFERENCES CITED, PTO-892.
- Note the attached INFORMATION DISCLOSURE CITATION, PTO-1449.

PART II.

A SHORTENED STATUTORY PERIOD FOR RESPONSE to comply with the requirements noted below is set to EXPIRE THREE MONTHS FROM THE "DATE MAILED" indicated on this form. Failure to timely comply will result in the ABANDONMENT of this application. Extensions of time may be obtained under the provisions of 37 CFR 1.136(a).

- Note the attached EXAMINER'S AMENDMENT or NOTICE OF INFORMAL APPLICATION, PTO-152, which discloses that the oath or declaration is deficient. A SUBSTITUTE OATH OR DECLARATION IS REQUIRED.
- APPLICANT MUST MAKE THE DRAWING CHANGES INDICATED BELOW IN THE MANNER SET FORTH ON THE REVERSE SIDE OF THIS PAPER.
 - a. Drawing informalities are indicated on the NOTICE RE PATENT DRAWINGS, PTO-948, attached hereto or to Paper No. 5 CORRECTION IS REQUIRED.
 - b. The proposed drawing correction filed on 5-7-92 have been approved by the examiner. CORRECTION IS REQUIRED.
 - c. Approved drawing corrections are described by the examiner in the attached EXAMINER'S AMENDMENT. CORRECTION IS REQUIRED.
- Formal drawings are now REQUIRED.

Any response to this letter should include in the upper right hand corner, the following information from the NOTICE OF ALLOWANCE AND ISSUE FEE DUE: ISSUE BATCH NUMBER, DATE OF THE NOTICE OF ALLOWANCE, AND SERIAL NUMBER.

Attachments:

- Examiner's Amendment
- Examiner Interview Summary Record, PTOL- 413
- Reasons for Allowance
- Notice of References Cited, PTO-892
- Information Disclosure Citation, PTO-1449
- Notice of Informal Application, PTO-152
- Notice re Patent Drawings, PTO-948
- Listing of Bonded Draftsmen
- Other

PARSHOTAM S. LALL
SUPERVISORY PATENT EXAMINER
ART UNIT 294



DS

**UNITED STATES DEPARTMENT OF COMMERCE
Patent and Trademark Office**

Address: Box ISSUE FEE
COMMISSIONER OF PATENTS AND TRADEMARKS
Washington, D.C. 20231

SAM PASTERNACK
CHOATE, HALL & STEWART
EXCHANGE PLACE
53 STATE STREET
BOSTON, MA 02109

**NOTICE OF ALLOWANCE
AND ISSUE FEE DUE**

Note attached communication from the Examiner
 This notice is issued in view of applicant's communication filed _____

SERIES CODE/SERIAL NO.	FILING DATE	TOTAL CLAIMS	EXAMINER AND GROUP ART UNIT	DATE MAILED
07/565,274	08/09/90	058	PIPALA, E	2304 06/30/92
First Named Applicant	DAVIS, JAMES R.			

TITLE OF INVENTION AUTOMOBILE NAVIGATION SYSTEM USING REAL TIME SPOKEN DRIVING INSTRUCTIONS (AS AMENDED)

ATTY'S DOCKET NO.	CLASS-SUBCLASS	BATCH NO.	APPLN. TYPE	SMALL ENTITY	FEES DUE	DATE DUE
2	364-443,000	E22	UTILITY	NO	\$1130.00	09/30/92

**THE APPLICATION IDENTIFIED ABOVE HAS BEEN EXAMINED AND IS ALLOWED FOR ISSUANCE AS A PATENT.
PROSECUTION ON THE MERITS IS CLOSED.**

THE ISSUE FEE MUST BE PAID WITHIN THREE MONTHS FROM THE MAILING DATE OF THIS NOTICE OR THIS APPLICATION SHALL BE REGARDED AS ABANDONED. THIS STATUTORY PERIOD CANNOT BE EXTENDED.

HOW TO RESPOND TO THIS NOTICE:

I. Review the SMALL ENTITY Status shown above.

If the SMALL ENTITY is shown as YES, verify your current SMALL ENTITY status:

- A. If the Status is changed, pay twice the amount of the FEE DUE shown above and notify the Patent and Trademark Office of the change in status, or
- B. If the Status is the same, pay the FEE DUE shown above.

II. Part B of this notice should be completed and returned to the Patent and Trademark Office (PTO) with your ISSUE FEE. Even if the ISSUE FEE has already been paid by a charge to deposit account, Part B should be completed and returned. If you are charging the ISSUE FEE to your deposit account, Part C of this notice should also be completed and returned.

III. All communications regarding this application must give series code (or filing date), serial number and batch number. Please direct all communications prior to issuance to Box ISSUE FEE unless advised to the contrary.

If the SMALL ENTITY is shown as NO:

- A. Pay FEE DUE shown above, or
- B. File verified statement of Small Entity Status before, or with, payment of 1/2 the FEE DUE shown above.

IMPORTANT REMINDER: Patents issuing on applications filed on or after Dec. 12, 1980 may require payment of maintenance fees.

0809

EXHIBIT

3

**Back Seat Driver: voice assisted automobile
navigation**

by

James Raymond Davis

B.S.A.D., Massachusetts Institute of Technology (1977)

Submitted to the Media Arts and Sciences Section
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1989

©Massachusetts Institute of Technology 1989
All Rights Reserved

Signature of Author

Media Arts and Sciences Section

August 4, 1989

Certified by *3-3-*

Nicholas P. Negroponte
Professor of Media Technology
Thesis Supervisor

Accepted by

Stephen A. Benton
Chairman, Departmental Committee on Graduate Students

The remainder of this document:

RITTMUELLER 187-351

matches (page-for-page and line-for-line) the document found at:

Ex. 2, pgs. 111-276.

To avoid unnecessary duplication, Harman will provide RITTMUELLER 187-351 at the Court's request.

EXHIBIT

4

Bellcore

© Bell Communications Research

I need a copy
of this sent to
Karen Lochbaum
Aiken Computation Lab
33 Oxford St
Cambridge MA
02138

Also enclosed
How to tell people ..
article in package

Lynn A. Streeter

445 South Street
Room MRE 2A-383
Box 1910
Morristown, New Jersey 07960-1910
201 829 4311
lynn@bellcore.com

STREETER 00052

**Back Seat Driver: voice assisted automobile
navigation**

by

James Raymond Davis

B.S.A.D., Massachusetts Institute of Technology (1977)

**Submitted to the Media Arts and Sciences Section
in partial fulfillment of the requirements for the degree of**

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1989

©Massachusetts Institute of Technology 1989

All Rights Reserved

Signature of Author

Media Arts and Sciences Section

August 4, 1989

Certified by

**Nicholas P. Negroponte
Professor of Media Technology
Thesis Supervisor**

Accepted by

**Stephen A. Benton
Chairman, Departmental Committee on Graduate Students**

The remainder of this document:

STREETER 52-217

matches (page-for-page and line-for-line) the document found at:

Ex. 2, pgs. 111-276.

To avoid unnecessary duplication, Harman will provide STREETER 52-217 at the Court's request.

EXHIBIT

5

Thesis Defense:
Telling You Where To Go
Jim Davis
Friday 26 May 1989
4 PM E15-283a

This thesis is about the design and construction of a machine that does something difficult in a new way. The task is to help people ~~people~~ find their way by car from one place to another within a city. This task is both useful and needy of improvement. The United States federal Highway Administration estimates that 45 billion dollars are wasted each year in the U.S. because of ineffective routing, from causes including being lost, stuck in traffic, or choosing bad routes.

The solution proposed here is called the Back Seat Driver. This name has two connotations: first, that of an unwanted critic of one's driving skills, the second that of a helpful agent who can direct one through a locale. It is the latter sense I intend. The Back Seat Driver is a computer program which uses synthetic speech to give instructions to the driver of a car as needed while driving. Until now, people needing directions have used their eyes to read instructions or maps (either paper or electronic). Using speech rather than text or graphics leaves the driver's eyes and hands free for the driving task.

I'll talk about how people give directions and how to make a program give directions that are just as good, or even better.

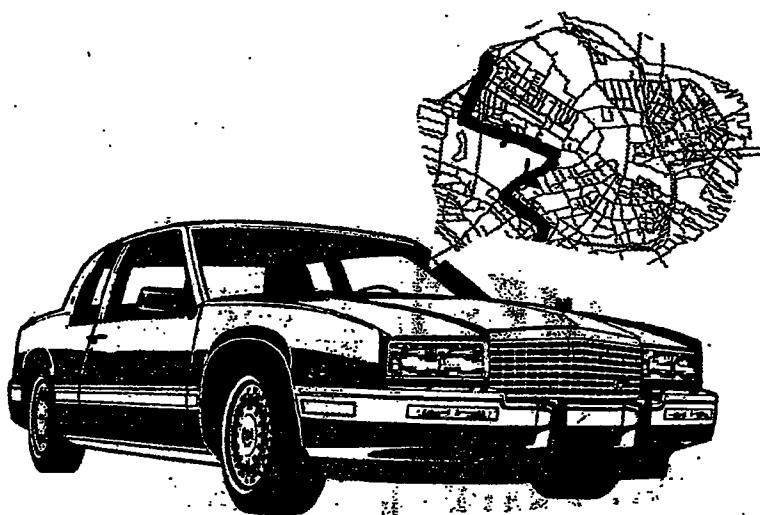


EXHIBIT 6

**FILED
UNDER SEAL**

EXHIBIT 7

**FILED
UNDER SEAL**

EXHIBIT 8

**FILED
UNDER SEAL**

EXHIBIT 9

**FILED
UNDER SEAL**

EXHIBIT

10

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MASSACHUSETTS**

**MASSACHUSETTS INSTITUTE OF
TECHNOLOGY,**

Plaintiff,

v.

**Case No: 05-10990 DPW
Hon. Douglas P. Woodlock**

**HARMAN INTERNATIONAL
INDUSTRIES, INCORPORATED,**

Defendant.

**MIT'S SUPPLEMENTAL RESPONSES TO
HARMAN'S INTERROGATORY NOS. 1-5, 8, 11, AND 13-14**

Pursuant to Rules 26 and 33 of the Federal Rules of Civil Procedure, Plaintiff, Massachusetts Institute of Technology ("MIT") submits the following supplemental responses and objections to Harman International Industries, Incorporated's ("Harman's") Interrogatory Nos. 1-5, 8, 11, 13, and 14 (the "Interrogatories").

GENERAL OBJECTIONS

MIT herein incorporates by reference its General Objections as set forth in MIT's First Supplemental Response to Harman's Interrogatory Nos. 11-15.

SPECIFIC OBJECTIONS AND RESPONSES

INTERROGATORY NO. 1

State MIT's proposed construction, and all bases supporting such construction, of the following element of claim 1 of U.S. Patent No. 5,177,685 (the "Patent-In-Suit"): "a map database functionally connected to said computing apparatus which distinguishes between physical and legal connectivity."

SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 1

Subject to, and without waiving, the foregoing general and specific objections, MIT incorporates by reference the preliminary claim constructions set forth in its May 17, 2006 correspondence to Harman. MIT asserts that “a map database functionally connected to said computing apparatus which distinguishes between physical connectivity and legal connectivity” should be construed as “a map database that contains information on both physical connectivity and legal connectivity and arranged so that the computing apparatus can gain access to this information.” MIT further asserts that “physical connectivity” should be construed as “how pieces of pavement connect or whether two segments touch.” MIT further asserts that “legal connectivity” should be construed as “whether one can legally drive onto a physically connected piece of pavement or whether it is legal to travel from one segment to another.” Support for these constructions can be found in the ‘685 patent at, for example, the abstract, FIGS. 1 & 3; col. 1, line 59-col. 2, line 33; col. 3, line 23-col. 4, line 10; col. 4, line 61-col. 5, line 25; and col. 7, line 48-col. 8, line 11.

MIT further states that additional bases and explanation may be discussed in MIT’s experts’ reports, and MIT intends to incorporate those reports herein by reference in their entirety in response to the Interrogatory.

MIT reserves the right to amend this Response after the Court issues its *Markman* order to the extent that the *Markman* order construes the claims differently than the construction asserted by MIT and used in responding to this Interrogatory.

INTERROGATORY NO. 2

For each asserted claim of the Patent-in-Suit, identify each element that MIT contends is written in means-plus-function form as permitted under 35 U.S.C. § 112 ¶6 and, for each such

EXHIBIT

10

element, identify the specified function for that element and the corresponding structure, material, or acts described in the specification as required by 35 U.S.C. § 112 ¶6. MIT's response to this interrogatory should appear in columns 2 and 3 of Table A in Appendix A and should include citations to all evidence (by column and line number, or by figure and item number) allegedly supporting MIT's identified function.

SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 2

Subject to, and without waiving, the foregoing general and specific objections, MIT incorporates by reference the preliminary claim constructions set forth in its May 17, 2006 correspondence to Harman. MIT further identifies the following limitations as means-plus-function limitations:

<p>Claim 1: driver input means functionally connected to said computing apparatus for entering data into said computing apparatus, said data including a desired destination,</p>	<p>Function: entering data into the computing apparatus, the data including a desired destination. Structures: a computer keyboard (or its functional equivalents); a cellular telephone keypad (or its functional equivalents); speech input (e.g., using speech recognition software); predictive spelling algorithms, or equivalent structures that allow the driver to input information to the computing apparatus. FIGS. 1, 2, & 4; col. 3, line 23-col. 4, line 10; col. 20, line 67-col. 21, line 13; col. 23, line 62-col. 24, line 60.</p>
<p>Claim 21: The automobile navigation system of claim 1 further comprising means for updating said map database.</p>	<p>Function: updating the map database. Structures: replacing a particular version or iteration of the map data with a new, more accurate version of the map data; or updating the map data by radio broadcast or equivalent structures that allow the map database to be updated. FIGS. 1 & 3; Col. 3, line 43-col. 4, line 10; col. 7, line 48-col. 8, line 11; col. 12, line 44-col. 13, line 21.</p>

<p>Claim 48: The automobile navigation system of claim 1 wherein said driver input means includes means for said driver to demand immediate instructions, or clarification or repetition of instructions already provided.</p>	<p>Function: allowing the driver to (a) demand immediate instructions, (b) demand clarification of instructions already provided, or (c) demand repetition of instructions already provided, or any combination thereof.</p> <p>Structures: a button or command (or functional equivalents thereof) on the driver input (including tactile inputs and voice inputs), or equivalent structures that allow the driver to demand immediate instructions, clarification of instructions already provided or repetition of instructions already provided.</p> <p>FIGS. 1, 2, & 4; col. 18, lines 32-41; col. 23, line 62-col. 24, line 60.</p>
<p>Claim 49: The automobile navigation system of claim 1 wherein said driver input means includes means for said driver to indicate to said automobile navigation system that a given instruction provided by said system is impossible to complete for some reason and that a new route must be calculated.</p>	<p>Function: to allow the driver to indicate that a given instruction is impossible to complete or to request that a new route be calculated.</p> <p>Structures: a button or command (or functional equivalents thereof) on the driver input (including tactile inputs and voice inputs); an error by the driver or by the system as determined by the system (or functional equivalents thereof); or equivalent structures that allow the driver of the automobile to provide data to the system that a particular action can no longer be taken or that a new route is preferred.</p> <p>FIGS. 1, 2, & 4; col. 10, line 24-col. 11, line 16; col. 19, lines 39-52.</p>
<p>Claim 50: The automobile navigation system of claim 1 wherein said driver input means comprises a voice recognition system to allow at least some driver input to be spoken.</p>	<p>Function: allowing the driver to provide data or information to the system via the driver's voice.</p> <p>Structures: voice recognition software (or its functional equivalents); voice recognition templates (or its functional equivalents); or equivalent structures that accept the driver's spoken output as an input and converts the output into commands for the computing apparatus that are used by the system to respond to the input.</p> <p>FIGS. 1, 2, & 4; col. 20, line 67-col. 21, line 13; col. 24, lines 45-60.</p>

MIT further states that additional bases and explanation may be discussed in MIT's experts' reports, and MIT intends to incorporate those reports herein by reference in their entirety in response to the Interrogatory.

MIT reserves the right to amend this Response after the Court issues its *Markman* order to the extent that the *Markman* order construes the claims differently than the construction asserted by MIT and used in responding to this Interrogatory.

INTERROGATORY NO. 3

For each element identified in response to Interrogatory No. 2 above, state whether MIT contends that Harman literally infringes that element (i) because the structure identical to that identified in response to Interrogatory No. 2 above for that element is found in each accused product and, if MIT contends that identical structure exists, provide a detailed explanation of where such identical structure for that element is found in each accused product; or because (ii) a structure equivalent to that identified in response to Interrogatory No. 2 above for that element is found in each accused product and, if MIT contends that equivalent structure exists, provide a detailed explanation of where such equivalent structure for that element is found in each accused product. MIT's response to this interrogatory should appear in columns 4 of Table A in Appendix A and should include citations to all evidence allegedly supporting MIT's contention and its identified structure.

SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 3

MIT objects to this Interrogatory to the extent that it misstates and mischaracterizes the legal standard for infringement. Subject to and without waiving the foregoing general and specific objections, MIT incorporates herein by reference its Infringement Contentions and documents cited therein previously served on Harman, the deposition transcript of Mr. Radomski (and Plaintiff's Exhibits 16-39), the deposition transcript of Mr. Montealegre (and Plaintiff's Exhibits 40-56), the deposition transcript of Mr. Jeske, Harman's Rule 30(b)(6) designee (and Plaintiff's Exhibits 104-124), and the June 12, 2006 Letter from John Pint to Craig Leavell.

All of Harman's U.S. Navigation Products literally infringe the claim elements identified in response to Interrogatory No. 2. Moreover, structures equivalent to those used by Harman's U.S. Navigation Products to facilitate entry of the driver's destination were known and/or used at least by the filing date of the '685 patent and do not, in any event, constitute after-arising equivalents.

Each of Harman's U.S. Navigation Products includes an interface (i.e., a human-machine interface or HMI) to a navigation core (i.e., a navigation computer) that performs navigation operations such as route-finding, map-matching, and route guidance (i.e., discourse generation). Each of Harman's U.S. Navigation Products includes an input to the HMI. The inputs include haptic inputs (such as a rotary control knob with a predictive spelling algorithm) or an audio input (using speech recognition algorithms). The inputs allow the driver to communicate with the HMI to enter a destination for the automobile. The HMI provides the destination to the navigation core. Harman's U.S. Navigation Products feature a "next letter tree" that is an algorithm that eliminates possible next letter choices as a driver enters the desired destination and is, thus, a predictive spelling algorithm.

Systems using haptic inputs allow the user to spell the destination by selecting letters on a display by manipulating the haptic input. For example, the user selects the first letter of the destination by rotating a rotary knob, rotary push-button, or directional pad until the correct letter appears on the display or corresponds to the letter to be selected. The user then depresses the rotary knob to select the letter. The haptic input satisfies the plain and ordinary definition of a keyboard, namely: "a hardware device consisting of a number of mechanical buttons (keys) which the user presses to input characters to a computer." Harman's U.S. Navigation Products using a haptic input include the following: the TrafficPro family of products (including for

aftermarket, Porsche, Ford, and Land Rover), the Online-Pro family (including for Porsche GT and Ferrari), Rb4, Crossfire, Harley Davidson, DVD Navimodule (for Audi), PCM 2 and PCM 2.1 (for Porsche), W211, W220, W221 (for Mercedes), and the Cascade (for Ferrari).

Mr. Jeske testified that all of Harman's U.S. Navigation Systems allow the driver to request the current driving recommendation (i.e., the next immediate instruction). The driver requests the current recommendation by one of the inputs to the HMI of Harman's U.S. Navigation Systems.

Mr. Jeske testified that all of Harman's U.S. Navigation Products recalculate a route if the driver deviates from the previously-calculated route. The driver, by the deviation, informs the navigation core that the previously-calculated route cannot be followed for some reason (e.g., obstruction or driver error).

All of Harman's U.S. Navigation Products allow the map database or the data therein to be updated. All of Harman's U.S. Navigation Products can use an updated version of the map data, by, for example, purchase and use of a compact disc with updated map data. Harman's U.S. Navigation Products may also include updates to the navigation core when new versions of the map database are used. Moreover, Harman's U.S. Navigation Products do not operate without a map database containing map data for the United States. The updated versions of map data must be compatible with the navigation core and other hardware components of Harman's U.S. Navigation Products. As such, each new version of the map database available for Harman's U.S. Navigation Products represent an update to the map database of Harman's system.

MIT further states that additional bases and explanation may be discussed in MIT's experts' reports, and MIT intends to incorporate those reports herein by reference in their entirety in response to the Interrogatory.

MIT reserves the right to amend this Response after the Court issues its *Markman* order to the extent that the *Markman* order construes the claims differently than the construction asserted by MIT and used in responding to this Interrogatory.

INTERROGATORY NO. 4

For each element identified in response to Interrogatory No. 2 above, state whether MIT contends that Harman infringes that element under the doctrine of equivalents and, if MIT contends that doctrine of equivalents infringement exists, provide a detailed explanation of where that element is found in each accused product and the bases for MIT's contention that the differences between the element and the accused product are insubstantial and that the element and the accused product perform substantially the same function in substantially the same way to achieve substantially the same result. MIT's response to this interrogatory should appear in column 6 of Table A in Appendix A and should include citations to all evidence allegedly supporting MIT's contentions.

SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 4

MIT objects to this Interrogatory to the extent that it calls for a legal conclusion as to literal infringement and infringement under the doctrine of equivalents. MIT further objects to this Interrogatory to the extent that it misstates the law and/or legal standard for the doctrine of equivalents.

MIT incorporates its Response to Interrogatory Nos. 2 and 3 in their entirety herein. To the extent that Harman contends that its U.S. Navigation Products do not literally infringe the

claims identified in response to Interrogatory No. 2, Harman's U.S. Navigation Products infringe the claims under the doctrine of equivalents.

All of Harman's U.S. Navigation Products include elements equivalent to the elements of the claims of the '685 patent identified in response to Interrogatory No. 2. All of Harman's U.S. Navigation Products include elements that perform the same function as the elements of the claims identified in response to Interrogatory No. 2 in the same way to achieve the same result. Moreover, the elements of Harman's U.S. Navigation Products for input of the driver's destination are insubstantially different from those elements described and claimed in the '685 patent.

The function disclosed by the "driver input means" of claim 1 is allowing a driver to provide a desired destination to the computing apparatus. The way to achieve the function is through an interface between the driver and the computing apparatus. The result is that the driver's desired destination is provided to the computing apparatus, and route calculation is facilitated based on the desired destination. As discussed above, all of Harman's U.S. Navigation Products feature an input (e.g., haptic or acoustic) that allow the driver to specify the desired destination. The interface communicates the driver's destination to the navigation core (i.e., through the HMI). The navigation core in Harman's U.S. Navigation Products uses the destination as an input to the route calculation algorithms used by Harman's systems.

MIT further states that additional bases and explanation may be discussed in MIT's experts' reports, and MIT intends to incorporate those reports herein by reference in their entirety in response to the Interrogatory.

MIT reserves the right to amend this Response after the Court issues its *Markman* order to the extent that the *Markman* order construes the claims differently than the construction asserted by MIT and used in responding to this Interrogatory.

INTERROGATORY NO. 5

For each element of each asserted claim not identified in response to Interrogatory No. 2 above, provide a detailed explanation of where that element is found in each accused product and a statement as to whether that element is alleged to be present literally or under the doctrine of equivalents. For each element that is alleged to be present under the doctrine of equivalents, MIT's explanation should include the bases for MIT's contention that the differences between the claim element and the structure in the accused product are insubstantial that the element and the accused product perform substantially the same function in substantially the same way to achieve substantially the same result. MIT's response to this interrogatory should appear in column 7 of Table A in Appendix A and should include citations to all evidence allegedly supporting MIT's contentions.

SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 5

MIT objects to this Interrogatory to the extent that it calls for a legal conclusion as to literal infringement and infringement under the doctrine of equivalents. MIT further objects to this Interrogatory to the extent that it misstates the law and/or legal standard for the doctrine of equivalents. MIT further objects to this Interrogatory as premature as discovery is ongoing in this case.

MIT incorporates its Response to Interrogatory Nos. 1-4 in their entirety herein. Subject to and without waiving the foregoing general and specific objections, MIT states that all of Harman's U.S. Navigation products literally infringe the Asserted Claims of the '685 patent as

identified in the 6/12/06 Letter from John Pint to Craig Leavell, namely: claims 1, 2, 7-9, 11-13, 19, 21, 23, 24, 27-29, 32, 34-36, 40, 41-46, 48, 49, 54, and 56 at least because an identical or equivalent structure to that claimed is present in each of Harman's U.S. Navigation Products.

All of Harman's U.S. Navigation Products, including the following: the TrafficPro family of products (including for aftermarket, Porsche, Ford, and Land Rover), the Online-Pro family (including for Porsche GT and Ferrari), Rb4, Crossfire, Harley Davidson, DVD Navimodule (for Audi), PCM 2 and PCM 2.1 (for Porsche), W211, W220, W221 (for Mercedes), and the Cascade (for Ferrari), read on the claims of the '685 patent because Harman's U.S. Navigation Products employ all of the same elements arranged in the same way as the claims of the '685 patent to provide navigation functionality to the driver of automobiles including Harman's U.S. Navigation Products.

To the extent that Harman contends that Harman's U.S. Navigation Products do not literally infringe the Asserted Claims, Harman's U.S. Navigation Products infringe such claims under the doctrine of equivalents at least because all of Harman's U.S. Navigation Products include an element equivalent to each element of the Asserted Claims. All of Harman's U.S. Navigation Products perform substantially the same function as each element of the Asserted Claims in substantially the same way to achieve substantially the same result. Moreover, all of Harman's U.S. Navigation Products include elements that are insubstantially different from each element of each of the Asserted Claims.

Support for MIT's position can be found at MIT's Infringement Contentions and documents cited therein previously served on Harman, the deposition transcript of Mr. Radomski (and Plaintiff's Exhibits 16-39), the deposition transcript of Mr. Montealegre (and Plaintiff's Exhibits 40-56), the deposition transcript of Mr. Jeske (and Plaintiff's Exhibits 104-124), and the

June 12, 2006 Letter from John Pint to Craig Leavell, all of which are incorporated herein by reference.

Additional bases and explanation may be discussed in MIT's experts' reports, and MIT intends to incorporate those reports herein by reference in their entirety in response to the Interrogatory.

MIT reserves the right to amend this Response after the Court issues its *Markman* order to the extent that the *Markman* order construes the claims differently than the construction asserted by MIT and used in responding to this Interrogatory.

INTERROGATORY NO. 8

Describe in detail the nature and total dollar amount of the "damages" that MIT has allegedly suffered due to Harman's alleged infringement, including a description of any "reasonable royalty" damages, the applicable royalty rate, the amount and type of Harman's sales or other activity to which the royalty rate should be applied, and the identification of all evidence that MIT contends supports any such alleged damages, rates, and amounts.

SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 8: HIGHLY CONFIDENTIAL

MIT objects to this Interrogatory as premature, because MIT has not yet completed its factual and legal analysis with regard to damages. MIT further objects to this Interrogatory because it seeks information protected by the attorney-client privilege, work product doctrine, and/or other applicable privileges or immunities.

Subject to and without waiving the foregoing general and specific objections, MIT responds as follows: MIT is entitled to all damages under the U.S. patent laws adequate to

compensate it for Harman's infringement of the '685 patent, together with interest and costs as fixed by the Court.

Harman has yet to produce documents or information detailing past sales of each accused product, but has provided a document bearing Bates number HAR278276, which Harman claims represents the total U.S. map database CDs/DVDs licensed for Harman's products. That document lists 321,260 sales through February 2006. Unless and until Harman produces better information, MIT will rely on that number as representing the total sales of the accused products through February 2006. The parties are currently in negotiations to obtain that further information.

MIT currently estimates that it will seek a reasonable royalty based on these past sales totaling at least between \$ 9,637,800 – \$11,244,100. On current information and belief, MIT will seek a per unit royalty amount for each past sale, which has not yet been determined, but which will be between \$ 30 and \$35, and will be set forth in MIT's expert's report, as will the bases and support for that amount.

MIT further contends that Harman's infringement is willful, and accordingly, MIT is entitled to increased damages up to three times the amount found or assessed, and costs and attorneys' fees.

MIT reserves the right to supplement this response seasonably as contemplated by the Federal Rules of Civil Procedure, and will provide further disclosure in the report of its expert with respect to damages.

INTERROGATORY NO. 11

The Schmandt and Davis, "Synthetic Speech for Real Time Direction-Giving" publication (MIT 01101-02) notes "field trials" of the Back Seat Driver that occurred more than

1 year before the filing date of U.S. Patent No. 5,177,685. Mr. Davis' thesis (*see* HAR 001479) also notes that the Back Seat Driver had been used more than 1 year before the filing date of U.S. Patent No. 5,177,685. For each asserted claim of United States Patent No. 5,177,685, identify each and every limitation of the claim that MIT contends was not embodied in a field trial prior to August 9, 1989, and explain in detail all bases for any contention by MIT that such field trials do not render each asserted claim of the '685 patent invalid under 35 U.S.C. § 102(b).

SECOND SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 11

MIT objects to this Interrogatory as overly broad, unduly burdensome, and not reasonably calculated to lead to the discovery of admissible evidence. MIT further objects to this Interrogatory because it calls for a legal conclusion with respect to validity. MIT further objects to this Interrogatory because it seeks information protected by the attorney-client privilege, work product doctrine, and/or other applicable privileges or immunities.

Subject to and without waiving the foregoing general and specific objections, MIT incorporates by reference its Supplemental Response to Interrogatory No. 11 as if fully set forth herein.

MIT further identifies pages 1-138 of MIT's 30(b)(6) deposition testimony as responsive to this Interrogatory.

INTERROGATORY NO. 13

For each claim of U.S. Patent No. 5,177,685, identify the date(s) on which the subject matter recited therein was first completely conceived, and identify by Bates number all documents or other material that evidence all such date(s) in any way.

THIRD SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 13

MIT objects to this Interrogatory as overly broad, unduly burdensome, and not reasonably calculated to lead to the discovery of admissible evidence. MIT further objects to this Interrogatory as premature to the extent that it calls for a legal conclusion with respect to conception. MIT further objects to this Interrogatory to the extent that it mischaracterizes the legal standard for conception. MIT further objects to this Interrogatory because it seeks information protected by the attorney-client privilege, work product doctrine, and/or other applicable privileges or immunities.

Subject to and without waiving the foregoing general and specific objections, MIT objects to the phrase "completely conceived", but states that the subject matter of the '685 patent was conceived before the filing date of the application on which the '685 patent issued. The details of the conception were fully described in answer to numerous questions to the inventors propounded during the deposition testimony of Dr. James R. Davis, Ph.D, Christopher M. Schmandt and MIT pursuant to Rule 30(b)(6) of the Federal Rules of Civil Procedure and in response to this Interrogatory, and those answers are herein incorporated by reference.

MIT further states that claims 1-4, 7-10, 14, 16, 19, 24, 27, 28, 42-44, 48, 55, and 57-58 were conceived at least as early as April 1988. MIT further states that claims 5, 6, 11-13, 15, 17, 18, 20-23, 25, 26, 29-41, 45-47, 49-54, and 56 were conceived at least as early as June of 1989.

MIT identifies documents bearing Bates numbers MIT00433-MIT00947, MIT01101-MIT01102, MIT01370-MIT01378, MIT01955-MIT02002, and MIT02155-MIT02274 as responsive to this Interrogatory. MIT further identifies pages 83-90, 104, 145-162, 179, and 287 of Mr. Schmandt's deposition transcript as responsive to this Interrogatory. MIT further identifies pages 79-91, 129, 167-169, 205, and 219-238 of Dr. Davis' deposition transcript as

responsive to this Interrogatory. MIT further identifies pages 1-138 of MIT's 30(b)(6) deposition testimony as responsive to this Interrogatory.

INTERROGATORY NO. 14

For each claim of U.S. Patent No. 5,177,685, identify the earliest date(s), if any, on which the subject matter recited therein was first actually reduced to practice, and identify by Bates number all documents or other material that supported all such date(s).

THIRD SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 14

MIT objects to this Interrogatory as premature to the extent that it calls for a legal conclusion with respect to reduction to practice. MIT further objects to this Interrogatory to the extent that it mischaracterizes the legal standard for reduction to practice. MIT further objects to this Interrogatory because it seeks information protected by the attorney-client privilege, work product doctrine, and/or other applicable privileges or immunities.

Subject to and without waiving the foregoing general and specific objections, MIT states that the subject matter of the '685 patent was reduced to practice before the filing date of the application on which the '685 patent issued. The details of the reduction to practice were fully described in answer to numerous questions to the inventors propounded during the deposition testimony of Dr. James R. Davis, Ph.D., Christopher M. Schmandt, and MIT under Rule 30(b)(6) of the Federal Rules of Civil Procedure, and in response to this Interrogatory, those answers are herein incorporated by reference.

MIT states that the following claims were reduced to practice at least as early as June of 1989: 1-4, 7-9, 21, 23-25, 32-37, 40, 41, 53, 55, 57, and 58. MIT further states that the following claims were reduced to practice at least as early as August 4, 1989: 10-12, 15, 19, 20, 26-31, 42-49, 51, 52, and 54. MIT further states that the following claims were reduced to practice at least

as early as the filing date of the '685 patent, August 9, 1990: 5, 6, 13, 14, 16-18, 20, 22, 38, 39, 50, and 56.

MIT identifies documents bearing Bates numbers MIT00433-MIT00947, MIT01101-MIT01102, MIT01370-MIT01378, MIT01955-MIT02002, and MIT02155-MIT02274 as responsive to this Interrogatory. MIT further identifies pages 83-90, 104, 145-162, 179, and 287 of Mr. Schmandt's deposition transcript as responsive to this Interrogatory. MIT further identifies pages 79-91, 129, 167-169, 205, and 219-238 of Dr. Davis' deposition transcript as responsive to this Interrogatory. MIT further identifies pages 1-138 of MIT's 30(b)(6) deposition testimony as responsive to this Interrogatory.

Dated: June 16, 2006

Respectfully submitted,

Massachusetts Institute of Technology,

By its Attorneys,



Steven M. Bauer (BBO# 542531)
Kimberly A. Mottley (BBO# 651190)
Jacob K. Baron (BBO # 652568)
John W. Pint (BBO# 660548)
PROSKAUER ROSE LLP
One International Place
Boston, Massachusetts 02110-2600
Phone: 617-526-9600
Fax: 617-526-9899

CERTIFICATION

I, the undersigned, have reviewed MIT's Supplemental Responses to Harman's Interrogatory Nos. 1-5, 8, 11, and 13-14. The responses set forth herein, subject to inadvertent or undiscovered errors or omissions, are based on and therefore necessarily limited by the records and information still in existence, presently recollected, thus far discovered in the course of preparation of the responses, and currently available to MIT. Consequently, MIT reserves the right to make any changes in or additions to any of these responses if it appears at any time that errors or omissions have been made therein or that more accurate or complete information has become available. Subject to the limitations set forth herein, said responses are true to the best of my present knowledge, information and belief.

I hereby certify under penalty of perjury that the foregoing is true and correct.

Executed on this ___th day of June, 2006.

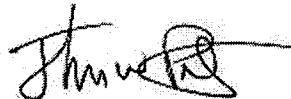
John H. Turner, Jr.
Associate Director, Technology Licensing Office
On behalf of Massachusetts Institute of Technology

CERTIFICATE OF SERVICE

I HEREBY CERTIFY that on June 16, 2006, I caused a true and correct copy of MIT's SUPPLEMENTAL RESPONSES TO HARMAN'S INTERROGATORY NOS. 1-5, 8, 11 AND 13-14 to be served on the following counsel of record via email:

Robert J. Muldoon, Jr.
SHERIN AND LODGEN, LLP
101 Federal Street
Boston, MA 02110

William A. Streff Jr., P.C.
Michelle A. H. Francis
Craig D. Leavell
Ann H. Chen
Colleen M. Garlington
KIRKLAND & ELLIS LLP
200 East Randolph Drive
Chicago, IL 60601
(312) 861-2000 (phone)
(312) 861-2200 (fax)



John W. Pint

EXHIBIT 11

**FILED
UNDER SEAL**

EXHIBIT 12

**FILED
UNDER SEAL**

EXHIBIT 13

**FILED
UNDER SEAL**

EXHIBIT 14

**FILED
UNDER SEAL**

EXHIBIT 15

**FILED
UNDER SEAL**

EXHIBIT 16

**FILED
UNDER SEAL**

EXHIBIT 17

**FILED
UNDER SEAL**

EXHIBIT 18

**FILED
UNDER SEAL**

EXHIBIT 19

**FILED
UNDER SEAL**

EXHIBIT

20

SYNTHETIC SPEECH FOR REAL TIME DIRECTION-GIVING

Christopher M. Schmandt and James Raymond Davis
The Media Laboratory
Massachusetts Institute of Technology

Abstract

The Back Seat Driver is a research prototype of a system to use speech synthesis as a navigational aid for an automobile equipped with localization equipment. We are evaluating the user interface by field trials. As this is work in progress, this paper will primarily give an overview of the system and describe its components. Included will be discussion of the map database, route finding algorithm, repair strategies, and the discourse generator.

With advances in navigation technology and automotive electronics[3,8] has come increasing interest in cars that know where they are and can help you figure out how to reach your destination. Most prototype projects have used various forms of display to present this information, and not all of them have included route finding ability[2,5,7,10,12,13,14,15,18,20,21]. For safety reasons, a display may not be particularly suited to this task, moreover there is some evidence that drivers do better following spoken directions than reading maps [19]. Our project, the Back Seat Driver, uses synthetic speech to give driving directions in real time. It plans a route, talks the driver through the route, and not only warns the driver when she has made an error, but also plans an alternate, corrective route.

This paper is an overview describing work in progress. We hope to publish more detailed explanations of the various portions at a later date. At the time of this writing (June, 1989) we have a working system on the road and are simultaneously conducting field trials and improving the direction giving ability and database. Although we do not aspire to prove that voice is better than graphics for direction giving, we do aim to build an optimal system. Early results are very encouraging, suggesting that speech may prove to be a powerful technology in automobiles of the future.

Talking about Directions

There are many factors which contribute to good route description by people, some of which our system only touches on. The problem is complex and simple solutions are not likely to produce comfortable interfaces.

A good route is not simply the shortest, but is more likely to be a combination of the fastest and the easiest to follow. "Easiest to follow" will, however, differ between directions given in advance and directions given in real time by a fellow passenger. Directions given in advance (as e.g. by [4], or the system at Hertz rental counters) must be simple, because the driver alone has the burden of interpreting and following the directions, and there is no help if the driver gets lost. When the direction giver is in the car it is practical to use minor streets or short cuts.

Good directions take into account conceptual portions of a route, which make it easier of the driver to keep track of her location on a more global basis. These may include named neighborhoods, types of neighborhoods (business, residential, parks) and types of roads (expressways, parkways, "main" roads, twisty or narrow streets).

By way of example, one of the authors was recently given directions at a car rental counter in a city new to him. The agent at the counter said "*As you leave the airport, keep bearing to the right. You'll go around the end of the runway and see signs for the Interstate north.*" The "computerized driving directions" printed at the counter described the same route as 5 separate segments, with mileages and names for each. Especially as it was night, the latter were almost impossible to follow, while the former had succinctly captured the salient aspects of the route.

When the directions are being given by a passenger, the real-time aspect becomes more important. Directions will be given just in time, taking into account vehicle speed, difficulty of the expected maneuver, driving styles, and road, weather, and traffic conditions. During long highway stretches with little need for description, the direction giver must maintain the driver's confidence. The passenger will also be watching for errors and trying to warn against them, again based on fine observations of the vehicle's speed and

direction. When a mistake is made, the passenger will tell the driver about it and together they will take corrective action (which is unlikely to be simply a sudden stop!).

Project Goals

The main goal of this project is to evaluate the utility of speech synthesis as the user interface to a real-time navigation system in an urban environment. Of particular concern is the discourse structure:

- how should driving acts be described?
- how can a description be generated from a route?
- how should timing considerations be applied?
- what kinds of feedback, both positive and negative, does the user require?
- what kinds of visual cues are most useful in describing an approaching location?

This information is gained from both laboratory simulations and field trials.

Our desire is to build the best possible real-time route describer. Although we believe a speech interface to the navigation unit is superior and safer than a visual interface, we do not plan to conduct direct comparison studies.

In the course of field trials to evaluate and improve our automatic direction giving, we hope to specify key components of the map database. We expect discourse behavior may need to vary with conditions (traffic, weather, day/night). It is likely that different visual cues may be useful in these situations. All must be represented in the database.

Geographic Database

Our database covers 41 square miles in the Boston area, including parts of Boston, Cambridge, Brookline, Somerville, and Watertown. It originated as a DIME (Dual Independent Map Encoding) file distributed by the United States Geological Survey[1]. A DIME file consists of a set of straight line segments, each with a name, a type, endpoints in longitude and latitude, and some additional information such as address numbers. Initially our database contained many errors. Correcting them required actually traveling most of the segments.

A DIME file alone is not sufficient for finding routes. The DIME files indicate physical connectivity, but route finding requires *legal* connectivity, i.e., one can legally drive from one segment to the next (one way streets are a simple example). We extended the data base format to explicitly

represent legal connectivity. Since some streets are better than others, the database must include a measure of quality. We take this to be a largely subjective measure of the ease of locating and following a street. This allows the route finder to prefer to use streets of higher quality.

The generation of easily followed natural descriptions required more extensions. We added a number of new segment types to distinguish bridges, underpasses, tunnels, rotaries, and access ramps. All these extensions were done for an earlier route finding project[4].

We are presently adding landmarks to the database. Drivers need landmarks to know how far to drive and when to turn. If the Back Seat Driver had eyes, it could simply choose landmarks as needed by looking for them in the landscape. Being blind, it must rely on landmarks coded into the map database. We have added traffic lights, stop signs, and some buildings to the portions of the landmark database. A main task now is to determine what else must be added.

In addition to landmarks, other information is useful for providing assistance following a route. We found it very useful to add lane information, both number of lanes as well as any turn restrictions on lanes (e.g. left turn only). On short street segments, it is important to give lane advice ("After the turn you'll want to get into the left hand lane.") or else the driver may be unable to make the following turn. Lane warnings ("Stay out of the left turn lane.") are also important driving cues.

An interesting problem arises at complex intersections, typically a maze of ramps between major arteries, possibly at different elevations (see figure 1). Such intersections are typically not accurately recorded in the map. Furthermore, limitations in the resolution of the position tracking equipment make it difficult to distinguish one segment from another, especially as they are likely to diverge at narrow angles. The combination of uncertainties in the map and uncertainties in position make it difficult to give a clear spoken directions. Fortunately these intersections are usually well signed, so the Back Seat Driver can give directions by referring to the signs, e.g. "Follow the signs to the expressways and airport". The wording of these signs needs to be in the database. It is important that Back Seat Driver's understand what the sign says, not simply utter the words. There are two reasons for this. First, our internal representation for text is based on syntactic structure, not text strings. Second, the objects mentioned in the signs (cities and roads) should be entered into the discourse model. They should become salient for future reference. This means that the text of a sign must be parsed, so that e.g. the sign text "Cambridge, Somerville, and Storrow Drive" should become a conjunction of the two cities "Cambridge" and "Somerville" and the street named "Storrow Drive".

EXHIBIT

20

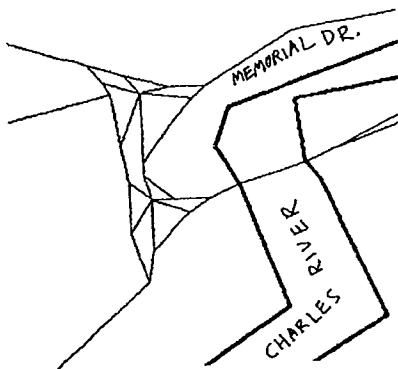


Figure 1: Access ramps at an interchange

System

Our vehicle is equipped with a localization unit built by NEC Home Electronics, Ltd., the project sponsor. It is a dead-reckoning position keeping system which uses speed and direction sensors. To compensate for error, it uses map matching on a map database stored on CD ROM. The system described more fully in [16,17]

As this is a research prototype, much of the computation is done in a base station computer laboratory (on a Symbolics Lisp Machine), rather than a computer on the vehicle. Two cellular telephones link the computer to the car. The on board navigational hardware transmits position and velocity via modem and cellular phone to the base station. The base station computer does all route planning and discourse generation. Speech synthesis is performed in a commercial text-to-speech synthesizer (Dectalk) cabled to the Lisp Machine. Synthesized instructions to the driver are relayed via the second cellular link and a speaker phone in the car. The keypad of the second phone also serves as the driver's control unit for the Back Seat Driver. Through this keypad a driver selects a destination, requests repeats of spoken information, and accesses other services of the Back Seat Driver.

A block diagram of the system appears in figure 2.

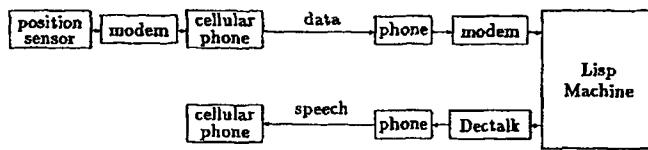


Figure 2: Communications block diagram

Routes

The Back Seat Driver can find the shortest, fastest, or most easily followed route. Route finding uses an A* search algorithm[11]. Depending on the driver's preference, one of three cost metrics is used. The *distance* metric is simply the sum of the lengths of the segments which compose a route. The *speed* metric scales each distance by a factor dependent on street quality. In addition, penalties are incurred for turns (more for left turns than right), stop signs, and traffic lights (which may be red). The *simplicity* metric, following [6], seeks to minimize the number of turns by imposing a distance penalty for each turn.

Discourse Strategies

The instructions are detailed and natural, and include a rich taxonomy of driving verbs. The dialog system uses cues such as vehicle speed and difficulty of driving actions to attempt to deliver instructions at the proper pace and in a timely manner. In addition, the system can anticipate some of the driver's possible mistakes and give warnings to avoid them.

Describing a route requires going from a series of segments (typically city blocks) in the database to a series of travel segments which will be separated by decision points. For example, going straight down a main street for five blocks will not be thought of by the driver as five separate acts, but rather one street traversal. A key piece of this analysis is that the need to make a decision is based on knowledge of what is *obvious*. Drivers do not want to be nagged at each corner to continue straight, but when they come to a questionable fork in the road they do want to be told which way to proceed.

If the driver does make a wrong turn, or misses a turn, the Back Seat Driver describes the error and then incrementally calculates a new route, rather than simply backtracking to the point of the error. Route planning includes weighting for length of the trip, difficulty of driving maneuvers (such as left turns against traffic), street quality, and complexity of the spoken directions.

As opposed to much prior work in discourse generation, the Back Seat Driver is a real-time system which must factor in a number of temporal considerations. It needs to give each stage in the directions at just the right point, in terms of the time it takes to execute the driving maneuver as well as the speed of the vehicle approaching the intersection. For safety considerations, we would rather err on the side of giving the driver plenty of warning, but a cue given too far in advance may be misused (e.g., a turn taken at an earlier intersection). Additionally, the software must consider the length of time it will take to recite an utterance. It is better to miss a turn and plan a new route than start describing the turn at a time when it may be unsafe to execute it (i.e., already well into an intersection).

There are several reasons to give instructions before the act, if time permits. One is to allow the driver to hear the instructions several times, and the other is to allow time to prepare for some acts, e.g., turns from a multi-lane street. These advance notices are lower priority than the description of the act itself, according to an internal set of system goals. Thus, they can be presented if there is adequate time, but will be ignored if the vehicle is approaching the next decision point too quickly.

Reassuring

While the driver is following a route, the system adopts a persistent goal of keeping the user reassured about her progress and the system's reliability. If Back Seat Driver were a human, this might be unnecessary, since the driver could see for herself whether the navigator was awake and attending to the road and driver. But the driver can not see the system, and so needs some periodic evidence that the system is still there.

One piece of evidence is the safety warnings the system gives (e.g. "slow down" before a turn), but if all is going well, there will not be any. The system gives two other kinds of evidence that things are going well. First, when the user completes an action, the system acknowledges the driver's correct action, saying something like "nice work" or "good". This feature is very popular with most test drivers.

The second form of evidence is to make insignificant remarks about the roads nearby, the weather, and so on. If the driver assumes that the navigator is being cooperative, as set out in Grice's maxims of cooperative conversation [9], then the driver can infer that everything is going well, for otherwise the navigator would not speak of trivial matters. It isn't clear, however, that one really wants a chatty speech synthesizer. Certainly this feature could be useful in a rented car in a new city, where it might actually have some interesting things to say.

Summary

The following is one of the more complex utterances of the Back Seat Driver to date. It summarizes many key points mentioned in this paper, and indicates the current state of operability of the discourse generator:

Get in the left lane because you're going to take a left at the next set of lights. It's a complicated intersection because there are two streets on the left. You want the sharper of the two. It's also the better of them. After the turn, get into the right lane.

The Back Seat Driver is already working in prototype form. Our present concerns are to determine what a spoken driving assistant should say, to understand how time and speed affect this decision, and to learn what features a map database must have to support generation of instructions.

Acknowledgments

The authors wish to gratefully acknowledge the support of NEC Home Electronics, Ltd.

References

- [1] *Geographic Base File GBDF/DIME: 1980 Technical Documentation*. U.S. Department of Commerce, Data Users Services Division, 1980.
- [2] Peter Braegas. Function, Equipment, and Field Testing of a Route Guidance and Information System for Drivers (ALI). *IEEE Transactions on Vehicular Technology*, VT-29(2):216-225, May 1980.
- [3] Donald F. Cooke. Vehicle Navigation Appliances. In *AUTO CARTO 7 International Symposium on Automation in Cartography*, pages 108-115, March 1985.
- [4] James R. Davis and Thomas F. Trobaugh. *Direction Assistance*. Technical Report 1, MIT Media Laboratory Speech Group, Dec 1987.
- [5] Ronald A. Dork. Satellite Navigation Systems for Land Vehicles. In *IEEE Position and Location Symposium*, pages 414-418, 1986. IEEE 86CH2365-5.
- [6] R. J. Elliot and M. E. Lesk. Route finding in street maps by computers and people. In *Proceedings of the National Conference on Artificial Intelligence*, pages 258-261, 1982.
- [7] Edward J. Krakiwsky et al. Research into electronic maps and automatic vehicle location. In *AUTO CARTO 8 International Symposium on Automation in Cartography*, 1987.

MIT 00936

- [8] Robert L. French. Automobile Navigation: Where is it going? In *IEEE Position and Location Symposium*, pages 406–413, 1986. IEEE 86CH2365-5.
- [9] H. P. Grice. Logic and conversation. In Cole and Morgan, editors, *Syntax and Semantics: Speech Acts*, pages 41–58, Academic Press, 1975.
- [10] Peter Haeussermann. *On Board Computer System for Navigation, Orientation, and Route Optimization*. Technical Paper Series 840483, Society of Automotive Engineers, 1984.
- [11] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on SSC*, 4:100–107, 1968.
- [12] Stanley K. Honey, Marvin S. White Jr., and Walter B. Zavoli. Extending Low Cost Land Navigation Into Systems Information Distribution and Control. In *IEEE Position and Location Symposium*, pages 439–444, 1986. IEEE 86CH2365-5.
- [13] Toshiyuki Itoh, Yasuhiko Okada, Akira Endoh, and Kenji Suzuki. *Navigation Systems Using GPS for Vehicles*. Technical Paper Series 861360, Society of Automotive Engineers, 1986.
- [14] M. D. Kotzin and A. P. van den Heuvel. Dead Reckoning Vehicle Location using a solid state rate gyro. In *Proceedings of the 31st IEEE Vehicular Technology Conference*, pages 169–172, April 1981. IEEE publication 81CH1638-6.
- [15] Ernst Peter Neukirchner and Wolf Zechnall. Digital Map Data Bases for Autonomous Vehicle Navigation Systems. In *IEEE Position and Location Symposium*, pages 320–324, 1986. IEEE 86CH2365-5.
- [16] Osamu Ono, Hidemi Ooe, and Masahiro Sakamoto. CD-ROM Assisted Navigation System. In *Digest of Technical Papers*, pages 118–119, IEEE ICCE, 1988.
- [17] Osamu Ono, Hidemi Ooe, and Masahiro Sakamoto. Navigation and communication system. In *Digest of Technical Papers*, IEEE ICCE, 1989.
- [18] Otmar Pilsak. *EVA: An electronic Traffic Pilot for Motorists*. Technical Paper Series 860346, Society of Automotive Engineers, 1986.
- [19] Lynn A. Streeter, Diane Vitello, and Susan A. Wonsiewicz. How to tell people where to go: comparing navigational aids. *International Journal of Man/Machine Systems*, 22(5):549–562, May 1985.
- [20] Katsutoshi Tagami et al. *Electro Gyro-Cator: New Intertial Navigation System for Use in Automobiles*. Technical Paper Series 830859, Society of Automotive Engineers, 1983.
- [21] Walter B. Zavoli and Stanley K. Honey. Map Matching Augmented Dead Reckoning. In *Proceedings of the 35th IEEE Vehicular Technology Conference*, pages 359–362, 1986. IEEE CH2308-5.

MR. SCHMANDT received his B.S. in Computer Science from MIT and an M.S. in computer graphics from MIT's Architecture Machine Group. He is currently a Principal Research Scientist and director of the Speech Research Group of the Media Laboratory at M.I.T.

His research interests are focused on interactive computer systems and human-interface issues of synchronous and asynchronous communication. His work emphasizes voice interaction for telecommunication based applications, with a goal of describing and then emulating human conversational behavior.



MIT 00937

The Back Seat Driver: Real Time Spoken Driving Instructions

James Raymond Davis and Christopher M. Schmandt
 The Media Laboratory
 Massachusetts Institute of Technology
 Voice: (617)-253-0314
 Fax: (617)-258-6264

Abstract

The Back Seat Driver is an automobile navigation aid which uses synthetic speech to give driving instructions in real time to the driver of a car. The advantage of speech over visual aids is that it leaves the driver's eyes free for driving, however it also poses special problems. This paper describes the strategies employed by the Back Seat Driver to successfully use speech. We hope this paper will persuade you of the value of speech in driving directions.

Introduction

The Back Seat Driver uses synthetic speech to give driving instructions in real time to the driver of a car. Speech is the only output channel it uses. There are no graphics. This paper discusses the advantages and problems arising from our exclusive use of speech to provide directions. The first section presents a brief overview of the Back Seat Driver. The second section describes the linguistic abilities of the Back Seat Driver. The final section describes the problems we have encountered because of our exclusive use of speech, and how we have overcome them.

System Overview

The architecture of the Back Seat Driver is shown in figure 1. At the center of the Back Seat Driver is the map database. The street map represents two ways in which streets can be connected: *physical* connectivity means it is physically possible to drive from one segment to another, and *legal* connectivity means it is lawful to do so. Legal connectivity is obviously needed to find legal routes, and physical connectivity for correctly describing intersections. The street map also includes traffic lights, stop signs, the number of lanes, and the location of all gas stations. These features are useful for both route finding (since, e.g. fast routes should avoid traffic lights) and for descriptions. The location system (supplied by the project sponsor, NEC) determines the current position of the vehicle by dead reckoning and map matching. It is further described in [3]. The driver gives the Back Seat Driver a destination by entering an address on a keyboard.

146 CH2789-6/89/0000-0146\$01.00 © 1989 IEEE

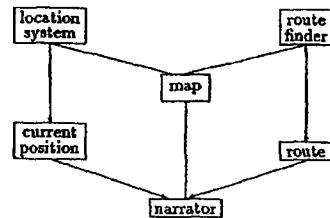


Figure 1: Back Seat Driver components

Using this map, the route finder can find the shortest route, the simplest one, or the one most easily followed, depending on the driver's preference.

The *narrator* is the subject of this paper. It generates instructions spoken by a speech synthesizer (a DecTalk). The narrator follows the driver's progress along the route. It decides what to say by comparing the current position against the map. The system follows the driver's progress, giving each instruction just when needed. If the time between instructions is long, the program gives the instruction twice, first in a detail, and later in a brief form. When not otherwise occupied, the system may deliver voice mail messages, weather reports, or commentary about the route. If the driver makes a mistake the system automatically finds an alternate route and continues.

The system has been running in prototype form since April 1989. It has been successfully used by drivers who have never driven in Boston. A somewhat longer description of the system appears in [4]. A complete description of the system appears in [1].

Linguistic Abilities

In designing the Back Seat Driver we chose to use speech as the sole means of providing driving instructions for two reasons. First, we believe that the driver's eyes are already employed watching traffic, and best left undisturbed. Second, we know that the alternative (map displays) will not work for those people who have difficulty reading maps[5]. We were also influenced by an experiment on route following which compared spoken instructions with paper maps[6]. Subjects who heard spoken directions did better than those with maps, and also better than those with both sources of guidance. Although this experiment does not compare real time speech to real time maps, it does suggest that spoken directions might be easier to follow than visual directions.

MIT 00938

Classifying Actions

Based on a study of how people naturally give spoken driving instructions, we developed a taxonomy of intersection types (Figure 2). This taxonomy is necessary in order to describe an intersection in the same way that a person would. For example, people talk about a "T" turn differently than a "fork" (or "Y") in the road. It is important that instructions match people's perceptions of the world they see.

The proper classification of an intersection depends upon the topology (how many streets are at an intersection), the geometry (the angles among them), and the types of roads involved. For instance, the difference between the "T" and "fork" mentioned above is one of geometry, not topology (figure 3), and the difference between a "fork" and an exit from a highway is that one of the two roads in the "Y" of the exit is much larger than the other.

In our system, a route is a sequence of street segments leading from the origin to the destination. We consider every connection from one segment to another as an "intersection", even if there is only one next segment at the intersection. At any moment, the car will be on one of the segments of the route, approaching an intersection (unless an error occurs, which is handled as discussed below). The task of the Back Seat Driver is to say whatever is necessary to get the driver to go from the current segment, across the intersection, to the next segment of the route.

The items in the taxonomy of intersection types are called *acts*. We use an object oriented programming methodology, so for each act there is a corresponding "expert". The Back Seat Driver generates speech by consulting these experts. At any moment, there will be exactly one expert in charge of telling the driver what to do. To select this expert, the Back Seat Driver asks each expert in turn to decide whether it applies to the intersection. The experts are consulted in a fixed order, the most specific ones first. The first expert to claim responsibility is selected. This expert then has the responsibility of deciding what (if anything) to say.

- CONTINUE
- FORCED-TURN
- TURN-AROUND
- TURN
- FORK
- ENTER
- EXIT
- ONTO-ROTARY
- EXIT-ROTARY
- STOP

Figure 2: Act taxonomy

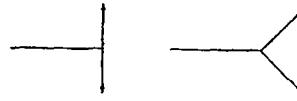


Figure 3: A "T" and a "Y" have the same topology

Describing actions

Each expert is able to generate text which describes the intersection. A description for an act must tell the driver two things: what to do and when (or where) to do it. "What to do" is expressed by a more or less constant verb phrase which depends upon the taxonomic classification, but may also depend upon specifics of the intersection. Thus a slight turn might be described by the verb "bear" where a sharper turn would be a "turn". The descriptions can be verbose or brief, and they can be expressed in past, present, or future tense. (We'll say why this flexibility is needed below.)

Saying "when"

Our study of natural instructions showed us that people almost never use distance as a cue for when to act. This is in sharp contrast to the textual directions provided by systems such as that of the Hertz rental company. Instead, people use two strategies. They wait until the driver is close to the intersection before saying anything, and/or they use a great variety of landmarks - including traffic lights, stop signs, other signs, buildings, road features, and the positions of other moving objects (e.g. "Follow that car."). The Back Seat Driver adopts both of these strategies.

Speech is especially useful as a cue for timing because speech is a temporal event, with a clear beginning and ending time. You know when someone begins to speak and when they finish. Someone peering at a map displayed on a CRT may have trouble distinguishing two adjacent streets, but there is no mistaking the word "now". Using time as a cue minimizes the workload on the driver, because the navigator absorbs the burden of remembering when to act. It also demands that the navigator have an accurate idea of where the car is. Our system demands positional accuracy of no greater than 10 meters for successful operation.

The Back Seat Driver's use of landmarks is unique in vehicle navigation systems. Our database began as a DIME file, but we extended it to include traffic lights, stop signs, road features (such as overpasses, bridges, and tunnels), distinctive signs, and the location of gas stations. Most of these are represented as attributes of the segments in the map database. To select a landmark for an intersection, the Back Seat Driver looks backwards from the intersection for the closest landmark which is also unique - that is, it prefers to say "take the first right after the underpass" rather than "take a right at the second set of lights". We think this makes the landmark easier to remember.

The Back Seat Driver does not speak at every single intersection. In the great majority of cases, it is perfectly obvious to the driver what to do (namely, to continue on forward). The action experts are also capable of deciding when the action at the intersection should be obvious to the driver. At present, the only action that is ever treated as obvious is CONTINUE. It is usually obvious to continue across an intersection, but we have found that what is obvious to one driver may not be so to another. Some people, for instance, are not comfortable driving across a major intersection unless they are instructed to do so. The expert can be somewhat customized so that its judgment of "obviousness" will correspond to that of the driver. If the action at the next intersection is obvious, the Back Seat Driver says nothing about it, and looks ahead for action at the next intersection, until it finds one that is not obvious.

The Back Seat Driver gives instructions just prior to the action. It also gives instructions further in advance, if time permits. This is especially useful when the instructions are complicated, as they are at some intersections. It is also able to give instructions "on demand". We call this the "what now" button. Drivers use this button for two reasons. Sometimes they are unsure whether they have come to the place where they are supposed to act, so they press the button to find out. At other times, they reach an intersection where the Back Seat Driver says nothing, because it believes the action is obvious, but it is not obvious to the driver. When the driver hits the "what now" button, the expert for the upcoming intersection describes it, even if it is considered to be obvious.

Talking about past and future

An advantage of language over pictures or gestures is that it can express events in the past or future. This advantage is well appreciated by readers of fiction, but may not yet be appreciated by designers of navigation systems. A navigation system should be able to talk about the past and future of the route, not just the present.

Drivers often need advance notice to prepare for an action. An example is what we call lane advice, which tells the driver to get into, or stay out of, a given lane. Lane advice is common in natural directions, and is one of the most appreciated features of the Back Seat Driver.

One reason for talking about the past is to describe mistakes. Drivers do not always follow the route the Back Seat Driver intends, either because of a mistake by the driver, the program, or external circumstances. When a mistake occurs, the Back Seat Driver finds a new route from the current location to the destination, while the driver is still moving. It also describes the mistake, saying something like "Oops, I meant for you to go straight." We think it is important that the system tell the user that there has been a mistake (without casting any blame on the user) so that the user will come to better understand the system's style of instruction giving, and so that the user will remain confident in the system's understanding of the route. Talking about past and future actions is important in navigation. Speech seems to be the easiest way of doing this.

Example

As an example, here's a sample of the description of the left turn from Fulkerson Street to Main Street in Kendall Square, Cambridge.

Get in the left lane because you're going to take a left at the next set of lights. It's a complicated intersection because there are two streets on the left. You want the sharper of the two. It's also the better of them. After the turn, get into the right lane.

This description was generated by the TURN expert in verbose form. It begins with some lane advice, then specifies the next action and provides a landmark for the place. The turn is described, and the proper street is described by two independent cues, one geometric, and one qualitative. Finally, the text provides a second piece of advice for after the turn.

Summary

The speech interface of the Back Seat Driver provides instructions without requiring the driver to look away from the road. Using speech permits us to talk about the past and the future as well as the present, and to give more detailed descriptions of the act than are possible with maps. Furthermore, it allows us to specify timing with great precision. But speech is not without its problems. The next section will discuss them, and the steps we have made to overcome them.

Liabilities of Speech

The advantages of a spoken language interface, as described above, do not come without cost. First, there are problems common to any natural language interface: while it is not terribly difficult to make a rudimentary interface, language generation requires substantial programming effort to be fluent and natural. Language is complicated, and people have literally a lifetime of experience with it, and are sensitive to fine nuances. On the other hand, having made this effort, we can exploit these nuances to convey extra information.

A second problem is that a natural language interface is only useful to those who speak the language. In our experience, only a few non-native speakers have been able to understand the directions. Map displays have conventions of their own, but are more universal than natural language. We have also noticed that some driving terms used in the Boston area (e.g. "rotary") are not in the dialect of other English speakers. In our view, universality is not a prime concern. We believe that systems should be custom fit to the idiosyncrasies of their owners. The Back Seat Driver in your car should speak to you in the language and terms that are best for you as an individual, not you as a generic human.

The remainder of this section discusses problems specific to spoken natural language generation.

Speech takes time

As we said above, speech is inherently temporal. We take advantage of this when we use speech as a timing cue, but it also can be a difficulty. A real time spoken navigation system must plan its speech to ensure that it has enough time to say what it needs to say. If little time remains, it must say less (or speak more quickly), or ask the driver to slow down. We handle this problem by tracking the vehicle's position and velocity, and by modeling the time required to speak. The Back Seat Driver begins its speech at a time chosen to be early enough to allow the driver to hear the entire message, understand it, and react to it, before the point where action must be taken. The model of reaction time includes a constant for the driver's comprehension and a variable time which depends on the speed of the car, according to the maximum comfortable braking deceleration.

The temporal nature of speech also requires that the Back Seat Driver sometimes combine instructions into a single utterance. When uttering an instruction, the Back Seat Driver looks ahead for the next instruction. If it determines that the time between the end of the execution of the current instruction and the beginning of the next is too short to allow it to speak the next instruction, it combines that text into the current one.

The Back Seat Driver does more than just give directions. Among other things, it also reads electronic mail messages from our office, gives weather reports, and makes comments about the route and road. Because speech takes time, and because a spoken utterance is only useful if completely spoken, the Back Seat Driver must carefully allocate the right to speak among potential tasks. It is undesirable for one task's speech to interrupt another's.

Speech can be misunderstood

A liability of speech, and synthetic speech in particular, is that speech can be misunderstood. This is particularly a problem with street names, because there are constraints that can help a driver correct a partially misunderstood name. A driver hearing an utterance that sounds like "Tarn left" can guess that it is a corrupt form of "Turn left", but nothing can help the driver know what was intended by "Tarn Street". Directions should not use street names, because street name signs may be hard to see, misaligned, or simply missing. The importance of this fact became apparent when we observed one driver who consistently misunderstood names, but also did not realize that he had misunderstood. Furthermore, the strength of his faith in the name was so strong that he drove straight through intersections, despite being told to "take the next left". This is probably the right thing to do with human instructions, where names are usually correctly understood, but street counts (e.g. "the third right") are imprecise or simply wrong. Our directions are phrased to minimize the use of street names in instructions. A typical text is: "Take the second left. It's Franklin Street."

Speech is transient

Information presented by speech does not persist, except in short term memory. We have already mentioned this as a reason why instructions should be given as late as possible. Another consequence of the transience of speech is that the system must be able to repeat itself at anytime, since the driver may not always be able to hear the speech. Repetition in turn poses a challenge.

since, unlike a program which reads the newspaper aloud, a literal repetition may not be appropriate, since the situation changes over time. For instance, if asked to repeat "Take the third left", the system may instead say "Take the second left" if the car has crossed an intersection. The consequence for the implementation is that the system retains not its previous words, but rather the previous reason for speaking. When asked to repeat, it invokes the same function that produced the last utterance.

A second problem with the ephemeral quality of speech is that the driver has no evidence of the program's existence except when it is speaking. We consider it very important that the driver have continued confidence that the program is running correctly, is aware of the driver's position and progress, and is "seeing" the world in the same way the driver does. We have devoted substantial effort to maintaining the illusion of co-presence.

In the introduction to this section, we said that the nuances of language could be used to convey much information. Co-presence is an idea communicated more by nuance than by explicit statement. (People would laugh if the system said "I'm right here with you." It sounds like something a therapist would say.) One way we indicate co-presence through nuance is by using deictic pronouns. Deictics are words that "point" at something. In English, we have four deictic pronouns: "this", "that", "these", and "those". The first two are singular, the second plural. The difference between "this" and "that" (and "these" and "those") is that "this" refers to something close. We use this in referring to landmarks. When the landmark is close, we use the proximal form (e.g. "these lights"); when distant, we use a brief noun phrase (e.g. "the next set of lights"). This is important. When a driver is stopped 30 meters back from a stop light, it may be literally true to say "turn left at the next set of lights", but it will confuse the driver.

A second means of conveying co-presence is to acknowledge the driver's actions. After the driver carries out an instruction the system briefly acknowledges the act if there is time, and if the act was not so simple (e.g. continuing straight) as to need no acknowledgment. This acknowledgment is a short phrase like "Okay". Some drivers dislike acknowledgments, so they can be disabled, but most find the confirmation comforting. The timing of the acknowledgment does much to confirm the driver's sense that the program really knows where the car is. Another source of acknowledgment is the use of cue words in the instructions. It will often be the case that the route calls for the driver to do the same thing twice (e.g. make two left turns). The speech synthesizer we use has very consistent pronunciation, and drivers sometimes get the impression that the system is

repeating itself because it is in error (like a record skipping). The acknowledgments help to dispel this, but we also cause the text to include cue words such as "another". These indicate that the system is aware of its earlier speech and the driver's previous actions.

Yet another means of conveying co-presence is to make occasional remarks about the road and the route. These remarks indicate that the program is correctly oriented. As an example, when the road makes a sweeping bend to one side, the program speaks of this as if it were an instruction ("Follow the road as it bends to the right.") even though the driver has no choice in what to do. The program also warns the driver about potentially hazardous situations, such as the road changing from one-way to two-way, or a decrease in the number of lanes. As with acknowledgments, these warnings can be disabled if the driver dislikes them. Other remarks have less to do with the route. We justify these by the maxims of cooperative conversations formulated by philosopher H. P. Grice[2]. His maxim of QUANTITY (part 1) says: "Make your contribution as informative as is required." Grice explains that one can convey information by appearing to flout the maxim. In this case, a driver can reason as follow: "The program, like all cooperative agencies, obeys the maxim of quantity. Therefore, it is had something important to say, it would say it. The program said nothing of great significance, therefore there is nothing urgently requiring my attention. So everything is well." At present, our "Gricean" utterances are trivial observations about the weather, but we are re-designing them to convey useful information about the city.

Summary

A speech interface for giving driving instructions has several advantages over a graphics interface. There are problems with natural language interfaces in general, and speech in particular, but they can all be overcome. The result is an excellent aid for navigation.

Acknowledgments

The authors wish to gratefully acknowledge the support of NEC Home Electronics, Ltd.

References

- [1] James Raymond Davis. *Back Seat Driver: voice assisted automobile navigation*. PhD thesis, Massachusetts Institute of Technology, September 1989.
- [2] H. P. Grice. Logic and conversation. In Cole and Morgan, editors, *Syntax and Semantics: Speech Acts*, volume 3, pages 41-58. Academic Press, 1975.
- [3] Osamu Ono, Hidemi Ooe, and Masahiro Sakamoto. CD-ROM Assisted Navigation System. In *Digest of Technical Papers*, pages 118-119. IEEE ICCE, 1988.
- [4] Christopher M. Schmandt and James R. Davis. Synthetic speech for real time direction giving. *IEEE Transactions on Consumer Electronics*, page (to appear), 1989.
- [5] Lynn A. Streeter and Diane Vitello. A profile of drivers' map reading abilities. *Human Factors*, 28:223-239, 1986.
- [6] Lynn A. Streeter, Diane Vitello, and Susan A. Wonsiewicz. How to tell people where to go: comparing navigational aids. *International Journal of Man/Machine Systems*, 22(5):549-562, May 1985.

EXHIBIT

21

06/21/90 18:59 FAX 617 258 6264

MIT/MEDIA LAB

10002

FAM 19.2

SYNTHETIC SPEECH FOR REAL TIME DIRECTION-GIVING

Christopher M. Schmandt and James Raymond Davis
 The Media Laboratory
 Massachusetts Institute of Technology

Abstract

The Back Seat Driver is a research prototype of a system to use speech synthesis as a navigational aid for an automobile equipped with localization equipment. We are evaluating the user interface by field trials. As this is work in progress, this paper will primarily give an overview of the system and describe its components. Included will be discussion of the map database, route finding algorithm, repair strategies, and the discourse generator.

Goals

The main goal of this project is to evaluate the utility of speech synthesis as the user interface to a real-time navigation system in an urban environment. Of particular concern is the discourse structure:

- how should driving acts be described?
- how can a description be generated from a route?
- how should timing considerations be applied?
- what kinds of feedback, both positive and negative, does the user require?
- what kinds of visual cues are most useful in describing an approaching location?

This information is gained from both laboratory simulations and field trials.

Our desire is to build the best possible real-time route describer. Although we believe a speech interface to the navigation unit is superior and safer than a visual interface, we do not plan to conduct direct comparison studies.

In the course of field trials to evaluate and improve our automatic direction giving, we hope to specify key components of the map database. We expect discourse behavior may need to vary with conditions (traffic, weather, day/night). It is likely that different visual cues may be useful in these situations. All must be represented in the database.

Geographic Database

Our database covers 41 square miles in the Boston area, including parts of Boston, Cambridge, Brookline, Somerville, and Watertown. It originated as a DIME (Dual Independent Map Encoding) file distributed by the United States Geological Survey[1]. A DIME file consists of a set of straight line segments, each with a name, a type, endpoints in longitude and latitude, and some additional information such as address numbers. Initially our database contained many errors. Correcting them required actually traveling most of the segments.

A DIME file alone is not sufficient for finding routes. The DIME files indicate physical connectivity, but route finding requires legal connectivity, i.e., one can legally drive from one segment to the next (one way streets are a simple example). We extended the data base format to explicitly represent legal connectivity. Since some streets are better than others, the database must include a measure of quality. We take this to be a largely subjective measure of the ease of locating and following a street. This allows the route finder to prefer to use streets of higher quality.

The generation of easy followed natural descriptions requires more extensions. We added a number of new segment types to distinguish bridges, underpasses, tunnels, rotaries, and access ramps. All these extensions were done for an earlier route finding project[2].

We are presently adding landmarks to the database. Drivers need landmarks to know how far to drive and when to turn. If the Back Seat Driver had eyes, it could simply choose landmarks as needed by looking for them in the landscape. Being blind, it must rely on landmarks coded into the map database. We have already added traffic lights to the landmark database. A main task now is to determine what else must be added.

System

Our vehicle is equipped with a localization unit built by NEC Home Electronics, Ltd., the project sponsor. It is a dead-reckoning position keeping system which uses speed and direction sensors. To compensate for error, it uses map matching on a map database stored on CD ROM. The system described more fully in [3].

As this is a research prototype, much of the computation is done in a base station computer laboratory (on a Symbolics Lisp Machine), rather than a computer on the vehicle. Two cellular

06/21/90 19:00 FAX 617 258 6264

MIT/MEDIA LAB

40003

telephones link the computer to the car. The on board navigational hardware transmits position and velocity via modem and cellular phone to the base station. The base station computer does all route planning and discourse generation. Speech synthesis is performed in a commercial text-to-speech synthesizer (Dectalk) cabled to the Lisp Machine. Synthesized instructions to the driver are relayed via the second cellular link and a speaker phone in the car. The keypad of the second phone also serves as the driver's control unit for the Back Seat Driver. Through this keypad a driver selects a destination, requests repeats of spoken information, and accesses other services of the Back Seat Driver.

A block diagram of the system appears in Figure 2, below.

Discourse Strategies

The instructions are detailed and natural, and include a rich taxonomy of driving verbs. The dialog system uses cues such as vehicle speed and difficulty of driving actions to attempt to deliver instructions at the proper pace and in a timely manner. In addition, the system can anticipate some of the driver's possible mistakes and give warnings to avoid them.

If the driver does make a wrong turn, or misses a turn, the Back Seat Driver describes the error and then incrementally calculates a new route, rather than simply back-tracking to the point of the error. Route planning includes weighting for length of the trip, difficulty of driving maneuvers (such as left turns against traffic), street quality, and complexity of the spoken directions.

As opposed to much prior work in discourse generation, the Back Seat Driver is a real-time system which must factor in a number of temporal considerations. It needs to give each stage in the directions at just the right point, in terms of the time it takes to execute the driving maneuver as well as the speed of the vehicle approaching the intersection. For safety considerations, we would rather err on the side of giving the driver plenty of warning, but a cue given too far in advance may be misused (e.g., a turn taken at an earlier intersection). Additionally, the software must consider the length of time it will take to recite an utterance. It is better to miss a turn and plan a new route than start describing the turn at a time when it may be unsafe to execute it (i.e., already well into an intersection).

Summary

The Back Seat Driver is already working in prototype form. Our present concerns are to determine what a spoken driving assistant should say, to understand how time and speed affect this decision, and to learn what features a map database must have to support generation of instructions.

Acknowledgments

The authors wish to gratefully acknowledge the support of NEC Home Electronics, Ltd.

References

- [1] *Geographic Base File GBDF/DIME: 1980 Technical Documentation*. U.S. Department of Commerce, Data Users Services Division, 1980.
- [2] James R. Davis and Thomas F. Trobaugh. *Direction Assistance*. Technical Report 1, MIT Media Laboratory Speech Group, Dec 1987.
- [3] Osamu Ono, Hidemi Ooe, and Masahiro Sakamoto. CD-ROM Assisted Navigation System. In *Digest of Technical Papers*, pages 118-119, IEEE ICCE, 1988.



Figure 1: Map of Boston area

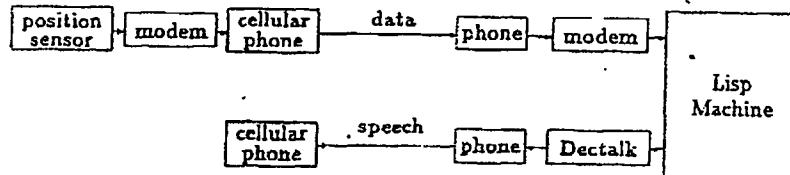


Figure 2: Block diagram of communications

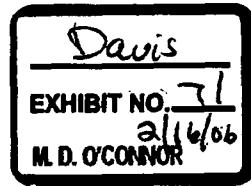
EXHIBIT

22

my life at the keyboard

My life at the keyboard

Jim Davis



Computers have been part of my culture all my life. My father worked for IBM. Every so often we'd visit him at work, and see the huge computer rooms. Sometimes he'd bring home pieces of computers for us to see. But I had no idea what these things really did. They were just cool looking. Sometime when I was 16 or 17 my father brought home some IBM documentation on programming languages and flowcharting. I tried to read them, but they did not make any sense to me - but this is not surprising, since IBM isn't exactly famous for clearly written documentation.

In 1973, I came to MIT. During the first few weeks, I went on a tour of the MIT Artificial Intelligence Laboratory, which in those days included the Logo Lab. I don't remember why I went on the tour, but I remember that I was with several other freshmen with whom I lived at my fraternity - but I can't recall if we were all together at the AI Lab because we were together as new members of the frat, or whether I came to join that frat because of liking the people I met at the AI Lab. The former seems more likely, especially since the house had a strong representation at the AI Lab. In any event, we were playing with Logo and were left pretty much alone - we had to figure it out by ourselves. We had a lot of fun with it.

So as an undergraduate I would sometimes come to the AI Lab to play (or "hack") with Logo, and I began to learn to program. I don't remember anyone teaching me - I think we must have taught each other. Sometimes I would try to use Lisp on the PDP-10 but it was too mysterious for me to figure out. That summer, I worked for IBM as an operator - a low-level position calling for about as much skill as an espresso maker. But in my spare time at work I was allowed to use the computer language APL, and this time I found a textbook for the language so I was able to learn some of it. I still had no idea how languages actually worked. I just used them.

In my sophomore year, I took the course 6.031 (which is the ancestor of the course now known as 6.001). This course explained how a computer language could be designed using a simpler language as building blocks. It also tried to give us some sense of the ideas of modularity and top down design, and most crucially, the idea of abstraction - that one can make a program which represents some concept or set of agreements, and thereafter use it without needing to know how the concept was implemented. The program becomes a "black box" whose internal details are irrelevant.

Later that year I took a second course which explained how the simplest sorts of computer languages (machine language) could be implemented by hardware circuits. I was now able to understand computer programming down to the level of individual "logic gates", if I wanted to. This reinforced my sense of the value of keeping different levels separate in order to build large, complex structure. Later, though, I would learn that one of the hardest problems is deciding where to draw the modularity lines, and that putting one's borders in the wrong place makes a system slow and difficult to use.

my life at the keyboard

The next major step in learning programming was a student job at the Architecture Machine Group (which is an ancestor of the Media Laboratory). In those days, a group of students at the ArcMac were developing a new operating system for use around the lab. The operating system, being new, was full of bugs, and these in turn demanded that there be constructed many software tools for examining the structures used by the program. I had the opportunity to look over the shoulders of those who were more experienced, and even to use the tools a bit to poke around. It was while using one of these tools that I suddenly understood that there is no actual meaning in the patterns of binary ones and zeros in machine, and no significant difference between the information on a disk and in the machine's memory. A given region of memory can be an instruction, or a number, or a letter, or a picture. The difference is solely a matter of interpretation. This was perhaps the biggest "aha" in my life, and I was happy that other people were around who could understand it and why it mattered.

It was while working at this same job that I began to think not just about how to make a program do something but how to make it easy for someone else to use. This was also a step towards being a professional programmer, a worker who makes artifacts for others to use, not just for his/her own delight in making it. It was also during this time that I also began to be good enough a designer that other people started taking my ideas for design and function seriously.

After I graduated, I began to work in the real world as a programmer. My first job was at Imlac in Needham MA. The Imlac was a minicomputer (what you'd call a workstation now) that was sort of an expanded PDP-8 with a built in vector display processor. It was programmed in assembly language. Imlac's big product was a phototypesetting system, CES, which took advantage of the raster graphics to offer a kind of WYSIWYG interface for the typesetting. This was before laser printers.

I kept this first job only a year, and then moved to a new job with the (Honeywell) Multics. Multics is an operating system of great historical importance. It was first developed as a partnership by MIT and General Electric as an experiment in a practical, very large time-sharing system. At the time, it was the very cutting edge of the state of the art in computer science. By the time I joined the Multics group, those days were past, but the group retained some measure of pride, and still had very high standards, even though time had passed them by. I learned several important ideas from working with the Multics people. First, my understanding of "interface" (the relation between a program and a user) expanded to include the idea that the user might be another programmer. It was important to make programs as building blocks by programmers whose needs you could not expect to easily anticipate. A second idea was that programs were meant to be read by people as well as machines. The Multics group had developed a programming process which required that all modification to the system be described and justified to a group of senior programmers before being written, and be read by some person other than the author before being installed. This was necessary because Multics was far too large for any single person to understand it. The coordination that this review board provided kept Multics stable and consistent as it grew and changed for more than 15 years. Though Multics is now nearly forgotten, it set a mark for software quality never equalled. Working with Multics taught me to be careful in my designs, to always to allow room for unanticipated future changes, and to expect people to read my programs.

my life at the keyboard

In my work at Multics I came to know many people, but one of particular note was Bernie Greenberg. Bernie was one of the most brilliant programmers I have ever met. In addition, he was a very talented musician, playing both rock guitar and baroque harpsichord with equal ease, and he spoke several languages. Bernie also re-introduced me to Lisp. At that time, the MIT Artificial Intelligence Laboratory was developing the first Lisp Machines and Bernie was friends with several of the key workers on this project. I learned Lisp from Bernie not long before he left Multics to join a new startup company to commercialize the Lisp Machine. I soon left as well, to join Logo Computer Systems, a new firm which intended to implement a version of Logo for the Apple II home computer.

Logo Computer Systems was the first time I was ever with a startup firm. Instead of the formal regulations of Multics, I was with an adhoc group which included several close friends and lovers, as well as some bizarre personalities. At LCSI we worked very, very hard, because we knew that money was in short supply. We would often work for 16 to 20 hours in a row. We did almost all our work in Lisp, on Lisp Machines, and I gradually became an expert with this language. In the end, we managed to produce our product on time, but then most of us left the company as a result of political battles with the higher management.

This turned out to be a blessing though, because Alan Kay had just gone to Atari, which was then quite rich, and Alan was setting up research labs in California and Cambridge. Almost the entire Boston staff of LCSI came to form the Atari Cambridge Research Center. Atari gave us money to design the best work environment we could think of, and freedom to work on problems that interested us. Not only was I able to work on music, I was able to hire one of my friends, Tom Trobaugh, to work with me. At Atari I knew the happiness of working with a partner on problems we really cared about using the most powerful computers available. Alas, Atari began to lose money, and one day it closed the lab.

After Atari went under I enrolled in the MIT's Media Lab, as one of the first contingent of doctoral students.

The Media Lab

The Media Lab encourages students to set their own directions, in fact it insists on it. This has both pros and cons. The advantage is that you learn to be independent, to think outside the common assumptions of the field. The drawback is that you don't always have the companionship of others while learning. In my own case, I became interested in the linguistic phenomenon called "paraverbals", those inarticulate noises like "uh huh" and "hmmm" that help make conversation run smoothly. There's a pretty large literature on the subject, but I had to discover it on my own, and I'm sure it would have gone faster with a guide. On the other hand, with an experienced authority controlling my learning, I wouldn't have done what I did.

The other important thing about the Media Lab is the constant focus on demonstrating one's work. Some people complain about it, but it's very important. To do a good demo, you have to be able to explain what you're doing and why it matters to a smart but uninformed person with just a few short sentences. You have to learn to be clear, and you have to learn to express your idea for the benefit of the learner, not the

my life at the keyboard

teacher. And you have to learn grace under pressure. I wish everyone learned these skills.

One of the first projects I worked on was the "phonetic dictionary". Remember when you were a kid, and you asked a grown-up for the meaning of a word, and they told you to look it up in the dictionary? That's a fine idea, except it's hard to do when you've only heard the word, and so don't know how it's spelled. This is especially true for English, with its bizarre spelling rules. The phonetic dictionary allows you to look up a word by spelling it according to how it sounds, not how it's spelled. You write some approximation of the sound of the word, and the system consults a dictionary that's organized by pronunciation. The key to the thing was being able to accept a wide range of "phonetic" spellings. For example, for the word "headache" you might write "hedayk" or "hedake". I was pleased to see this project mentioned on the very first page of Stewart Brand's book about the Lab.

My supervisor was Chris Schmandt, known to all by his login name "geek". I could write at length about his knowledge, but there are lots of smart people in the world. He has two qualities that are more rare. First, he's no autocrat. You can argue with him. There's no way to put on an air of superiority when you call yourself "geek". Second, and even more valuable, he kept his perspective. In particular, he was always taking off for a week or two at a time to bring his (then) baby daughter out into the wilderness. I learned from him that a computer will happily sit idle for a week, while a week lost from fathering is gone for ever.

My major project was the "Back Seat Driver", which was a car that could give you driving instructions in the city of Boston. It had a street map (so it knew the roads), a navigation system (so it knew where it was), and a speech synthesizer (so it could talk to you.). To actually make this work, I needed a car, and not just any car. The navigation system was supplied by our sponsor, a Japanese electronics firm, and was designed to work with only one type of car, a top of the line luxury sedan. I also needed not one but two cellular phones in the car for the communications. The Media Lab bought me what I needed and I kept the keys. I was surely the only graduate student in the USA with such lab equipment.

One day I was demonstrating the Back Seat Driver to a group from General Motors. When I took them out for a ride, they had a great time shooting pictures of each another getting into a Japanese car. Off we went driving. As usually happened, at one point the driver missed a turn. Normally, the consequence was that the BSD would calmly inform the driver of the fact and plan a new route, only in this case, the driver was a former race car driver, and he quickly made an (illegal) U Turn without even slowing down, a maneuver even seasoned Boston drivers never attempted. This caused the program to crash, but I guess that's better than crashing the car.

I'll have to add something here about the demise of Lisp Machines and the rise of Unix, and about becoming obsolete.

EXHIBIT 23

**FILED
UNDER SEAL**

EXHIBIT 24

**FILED
UNDER SEAL**

EXHIBIT

25

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MASSACHUSETTS**

**MASSACHUSETTS INSTITUTE OF
TECHNOLOGY,**

Plaintiff,

v.

**HARMAN INTERNATIONAL
INDUSTRIES, INCORPORATED,
A Delaware Corporation,
Defendant.**

**Case No: 05-10990 DPW
Hon. Douglas P. Woodlock**

**MIT'S RESPONSES TO HARMAN'S FIRST SET OF REQUESTS
FOR THE PRODUCTION OF DOCUMENTS AND THINGS (Nos. 1-29)**

Pursuant to Rules 26 and 34 of the Federal Rules of Civil Procedure, Plaintiff, Massachusetts Institute of Technology ("MIT") submits the following responses and objections to Harman International Industries, Incorporated's ("Harman's") First Set of Requests for the Production of Documents and Things (Nos. 1-29) (the "Requests")¹:

GENERAL OBJECTIONS

1. MIT objects to the Requests to the extent that they seek documents not relevant to the subject matter of the present lawsuit and/or are not reasonably calculated to lead to the discovery of admissible evidence.
2. MIT objects to the Requests to the extent that they seek documents or materials that constitute or contain information protected by the attorney-client privilege, the work-product doctrine, and/or any other applicable privilege. Any disclosure of any document that is subject to such a privilege or protection is inadvertent and shall not constitute a waiver of any such

¹ Although Harman identified the Requests as including "Nos. 1-28" under the caption and in the Certificate of Service, Harman served twenty-nine Requests.

privilege or protection, and any such document should, if located by Harman, be promptly returned to MIT. Any information or documents withheld on these grounds will be identified on a separate privilege log, except for such documents that were created or that came into being after initiation of this lawsuit.

3. MIT objects generally to the Requests to the extent they seek confidential documents or materials. MIT will produce confidential information only subject to an agreed-upon and appropriate protective order entered by the Court that limits the use of documents and information produced by MIT in this action to this proceeding and protects any documents containing MIT confidential information called for by Harman's Requests from unauthorized disclosure to Harman representatives or any third parties.

4. MIT objects generally to the Requests to the extent that they attempt or purport to impose obligations on MIT beyond those required by the Federal Rules of Civil Procedure 26 and 34, or the Local Rules of Practice of the United States District Court for the District of Massachusetts.

5. MIT objects to the Requests to the extent that they seek documents already in Harman's possession, equally accessible to Harman and/or publicly available.

6. MIT objects generally to the Requests to the extent that they seek documents or materials not in MIT's possession, custody, or control.

7. MIT objects generally to the Requests to the extent that they seek to characterize evidence in this matter. To the extent that MIT expressly or impliedly adopts any term used by the Requests, such adoption is specifically limited to these responses and does not constitute an admission of fact or law in this action or any other matter.

8. None of the General Objections and Responses herein is a direct or indirect admission (i) of the truth or accuracy of any statement or characterization asserted by Harman in any pleading or other filing with the Court; (ii) of the validity of any objection or response by Harman to any discovery request propounded in this action by MIT; or (iii) that any discovery request propounded by MIT in this action is wholly or partially objectionable under any applicable law or rules.

9. MIT generally objects to the Requests to the extent that MIT has not completed its investigation of the facts in this case and has not completed discovery in this action. Any responses to these Requests are based on information presently known to MIT and are without prejudice to MIT's right to supplement its responses and produce as evidence, any subsequently discovered documents.

10. MIT generally objects to the Requests to the extent they are overly broad or unduly burdensome, including without limitation Requests that require production of "all" or "any" documents.

11. MIT generally objects to the Requests to the extent they are vague, ambiguous, confusing, incomprehensible and/or unanswerable because of undefined or ill-defined terms and/or confusing syntax, or they fail to describe with reasonable particularity the documents sought.

12. MIT's willingness to provide documents or things is not a concession that the subject matter of the particular document is discoverable, relevant to this action, or admissible as evidence.

13. MIT expressly reserves all objections as to relevance, authenticity, and/or admissibility of any documents produced.

14. MIT has responded to the Requests as it interprets and understands each Request made therein. If Harman subsequently asserts an interpretation of any Request that differs from the understanding of MIT, MIT reserves the right to supplement its objections and responses.

15. The presence or absence of any general or specific objection does not mean that MIT does not object on any other grounds.

16. MIT incorporates its General Objections to Harman's First Set of Interrogatories (Nos. 1-7) as if fully set forth herein.

17. MIT objects to the production of any document falling within the General Objections set forth above. A statement by MIT that it will produce responsive, non-privileged, and non-work product protected documents means that MIT will produce documents, if any exist, within its possession, custody, or control that (a) are responsive to the particular Request or fall within the particular description in question, (b) have been located by MIT after a reasonable search, and (c) do not fall within one of the General Objections set forth above or within any of the objections contained in the specific responses set forth below.

18. Subject to these General Objections and the Specific Objections set forth below, MIT will produce responsive documents at a time and place mutually agreed upon by the parties.

SPECIFIC OBJECTIONS AND RESPONSES

REQUEST NO. 1

All documents and things referring or relating to U.S. Patent No. 5,177,685 (the "Patent-In-Suit"), the application that lead to its issuance (application number 565,274), or the subject matter thereof, including without limitation all engineering notebooks, laboratory notebooks, records, logs, and files created or maintained by, or at the direction of, Dr. James R. Davis, Mr. Christopher M. Schmandt or Dr. Sam Pasternack.

RESPONSE TO REQUEST NO. 1

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this request to the extent it seeks documents that are beyond the custody or control of MIT and are equally available to Harman. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

Subject to and without waiving the foregoing general and specific objections, MIT will produce responsive, non-privileged, and non-work product protected documents.

REQUEST NO. 2

All documents and things upon which MIT relied, at the time MIT filed its Complaint, to support the allegations that Harman has infringed the Patent-In-Suit.

RESPONSE TO REQUEST NO. 2

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine.

REQUEST NO. 3

All documents and things upon which MIT relied, at the time MIT sent its March 24, 2003 letter to Porsche (See Ex. A to Harman's Complaint in Case No. 05 C 1481), to support the assertion that Porsche needs "to license on a non-exclusive basis the right to practice these patents."

RESPONSE TO REQUEST NO. 3

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this request to the extent it inaccurately characterizes any statement made by or on behalf of MIT, in particular MIT's language in offering Porsche a license.

REQUEST NO. 4

All documents and things upon which MIT relied, at the time MIT sent letters to competitors or potential customers of Harman such as Garmin International, Inc., Alpine Electronics, Inc., Siemens AG, DaimlerChrysler AG, Bayerische Motoren Werke AG (BMW) and/or any other company to support the assertion that these companies need to license the Patent- In-Suit.

RESPONSE TO REQUEST NO. 4

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this request to the extent it inaccurately characterizes any statement made by or on behalf of MIT, in particular MIT's language in offering any entity a license.

REQUEST NO. 5

All documents and things referring or relating to Harman and/or any Harman product.

RESPONSE TO REQUEST NO. 5

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks

documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine.

Subject to and without waiving the foregoing general and specific objections, MIT will produce responsive, non-privileged, and non-work product protected documents related to the products in suit.

REQUEST NO. 6

All documents and things referring or relating to Back Seat Driver.

RESPONSE TO REQUEST NO. 6

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

Subject to and without waiving the foregoing general and specific objections, MIT will produce responsive, non-privileged, and non-work product protected documents.

REQUEST NO. 7

All documents and things referring or relating to any disclosure of any portion of the subject matter of the Patent-In-Suit to any third party.

RESPONSE TO REQUEST NO. 7

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT particularly objects to the Request to the extent it fails to

specify a relevant time period. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT and equally available to Harman. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

REQUEST NO. 8

All documents and things referring or relating to any disclosure to, or use by, any third party of any portion of the subject matter of the Patent-In-Suit before the filing date of patent application number 565,274.

RESPONSE TO REQUEST NO. 8

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT particularly objects to the Request to the extent it fails to specify a relevant time period. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

REQUEST NO. 9

All documents and things concerning the first sale, first offer for sale, or first public use or display of any alleged invention of the Patent- In-Suit.

RESPONSE TO REQUEST NO. 9

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to the Request as irrelevant to the extent it seeks documents relating to the sale of commercial products. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

REQUEST NO. 10

All documents and things referring or relating to any experimental use, field test, and/or prototype of any alleged invention of the Patent-In-Suit.

RESPONSE TO REQUEST NO. 10

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further

objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

REQUEST NO. 11

Documents sufficient to show each and every prototype, project, or product conceived, designed, researched, manufactured, used, sold, offered for sale, or imported into the United States allegedly within the scope of any claim of the Patent-In-Suit, including without limitation any activity of any third party in making, using, selling or offering for sale any product allegedly within the scope of any claim of the Patent- In-Suit.

RESPONSE TO REQUEST NO. 11

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege and/or the work product doctrine. MIT further objects to this Request as premature to the extent that it calls for legal conclusions with respect to claim construction and claim scope. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

REQUEST NO. 12

All documents and things referring or relating to Robert Swartz and/or Web Telephony, LLC.

EXHIBIT

25

RESPONSE TO REQUEST NO. 12

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to the Request to the extent it is irrelevant to the lawsuit. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT will produce the employment agreement between Mr. Swartz and MIT.

REQUEST NO. 13

All documents and things relating to any assertion, charge, claim, or allegation of infringement of the Patent-In-Suit, including without limitation correspondence, pleadings, discovery sought or obtained, transcripts, exhibits, videotapes, expert reports, or opinion.

RESPONSE TO REQUEST NO. 13

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it is unaware of any past or pending lawsuits, other than the present suit, alleging infringement of U.S. Patent No. 5,177,685.

REQUEST NO. 14

All documents and things referring or relating to MIT's decision to assert the Patent-In-Suit against Harman or any third party.

RESPONSE TO REQUEST NO. 14

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine.

REQUEST NO. 15

All documents and things evidencing, referring or relating to any communication between MIT, or anyone acting or purporting to act on MIT's behalf, and any party other than MIT, regarding the Patent-in-Suit or this litigation.

RESPONSE TO REQUEST NO. 15

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT will produce responsive, non-privileged, and non-work product protected documents.

REQUEST NO. 16

Documents sufficient to show the complete chain of title for, and MIT's ownership of and standing to sue under, the Patent-in-Suit.

RESPONSE TO REQUEST NO. 16

MIT objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT will produce responsive, non-privileged, and non-work product protected documents.

REQUEST NO. 17

All documents referring or relating to any prior art known, considered, reviewed or evaluated before or during prosecution, or after the grant of, the Patent-In-Suit.

RESPONSE TO REQUEST NO. 17

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

REQUEST NO. 18

All documents referring or relating to any search, investigation, study, analysis, evaluation or opinion relating to prior art, patentability, novelty, validity, enforceability, ownership and/or infringement of any claim of the Patent-In-Suit and the results of any such search or investigation.

RESPONSE TO REQUEST NO. 18

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

Subject to and without waiving the foregoing general and specific objections, MIT states that it is unaware of any non-privileged and/or non-work product protected documents responsive to this request, other than documents created by Harman or its representatives.

REQUEST NO. 19

Documents sufficient to show MIT's document retention policies from 1987 to present.

RESPONSE TO REQUEST NO. 19

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further

objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has not maintained a formal document retention policy at any time from 1987 to present.

REQUEST NO. 20

All documents and things referring or relating to MIT's first knowledge of any Harman product that MIT alleges infringes the Patent-In-Suit.

RESPONSE TO REQUEST NO. 20

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT will produce responsive, non-privileged, and non-work product protected documents.

REQUEST NO. 21

All documents and things passing between MIT and any named inventor of the Patent-In-Suit.

RESPONSE TO REQUEST NO. 21

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks

information that is confidential or is protected by attorney-client privilege and/or the work-product doctrine.

REQUEST NO. 22

Copies of all patents for which any named inventor of the Patent-In-Suit is named as an inventor.

RESPONSE TO REQUEST NO. 22

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

REQUEST NO. 23

The personnel, student and/or employment history files for each of the named inventors of the Patent-In-Suit and any MIT students and/or employees disclosed in MIT's initial disclosures, including curricula vitae.

RESPONSE TO REQUEST NO. 23

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT.

REQUEST NO. 24

All documents and things concerning any opinion concerning the validity, enforceability, or infringement by any person of the Patent-In-Suit.

RESPONSE TO REQUEST NO. 24

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

Subject to and without waiving the foregoing general and specific objections, MIT will produce responsive, non-privileged, and non-work product protected documents.

REQUEST NO. 25

All documents referring or relating to Charles G. Call, including without limitation all documents provided to or received from Charles C. Call.

RESPONSE TO REQUEST NO. 25

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT will produce responsive, non-privileged, and non-work product protected documents.

REQUEST NO. 26

All documents referring or relating to the funding or sponsorship for the Media Lab, including but not limited to the work and/or project that led to the Patent-In-Suit.

RESPONSE TO REQUEST NO. 26

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence.

REQUEST NO. 27

All thesis drafts, revisions, notes and documents referring or relating or leading up to the September 1989 thesis paper authored by James R. Davis, entitled Back Seat Driver: voice assisted automobile navigation.

RESPONSE TO REQUEST NO. 27

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks information protected by attorney-client privilege and/or the work-product doctrine. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

REQUEST NO. 28

Documents sufficient to show MIT's thesis submission, review, certification, acceptance and publication process(es).

RESPONSE TO REQUEST NO. 28

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT particularly objects to the Request to the extent it fails to specify a relevant time period. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request because validity and enforceability of the '685 Patent are not presently at issue in this case.

REQUEST NO. 29

Documents sufficient to show all incentives or rewards by MIT to inventors for issued patents.

RESPONSE TO REQUEST NO. 29

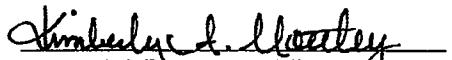
MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent that it is not relevant to the lawsuit.

Dated: October 11, 2005

Respectfully submitted,

Massachusetts Institute of Technology,

By its Attorneys,


Steven M. Bauer (BBO# 542531)
Kimberly A. Mottley (BBO# 691190)
PROSKAUER ROSE LLP
One International Place
Boston, Massachusetts 02110-2600
Phone: 617-526-9600
Fax: 617-526-9899

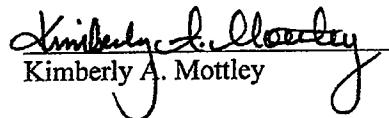
CERTIFICATE OF SERVICE

I HEREBY CERTIFY that on October 11, 2005, I caused a true and correct copy of
MIT'S RESPONSES TO HARMAN'S FIRST SET OF REQUESTS FOR THE PRODUCTION
OF DOCUMENTS AND THINGS (NOS. 1-29) to be served on the following counsel of record
via email, with a courtesy copy by U.S. Mail:

Robert J. Muldoon, Jr.
SHERIN AND LODGEN, LLP
101 Federal Street
Boston, MA 02110

William A. Streff Jr., P.C.
Michelle A. H. Francis
Jamal M. Edwards
KIRKLAND & ELLIS LLP
200 East Randolph Drive
Chicago, IL 60601
(312) 861-2000 (phone)
(312) 861-2200 (fax)

By:



Kimberly A. Mottley

EXHIBIT

26

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MASSACHUSETTS**

**MASSACHUSETTS INSTITUTE OF
TECHNOLOGY,**

Plaintiff,

v.

**HARMAN INTERNATIONAL
INDUSTRIES, INCORPORATED,**

Defendant.

**Case No: 05-10990 DPW
Hon. Douglas P. Woodlock**

**MIT'S RESPONSES TO HARMAN'S SECOND SET OF REQUESTS
FOR THE PRODUCTION OF DOCUMENTS AND THINGS (Nos. 30-61)**

Pursuant to Rules 26 and 34 of the Federal Rules of Civil Procedure, Plaintiff, Massachusetts Institute of Technology (“MIT”) submits the following responses and objections to Harman International Industries, Incorporated’s (“Harman’s”) Second Set of Requests for the Production of Documents and Things (Nos. 27-58) (the “Requests”)¹:

GENERAL OBJECTIONS

The following general statements and objections are incorporated into each of MIT’s responses, as set forth there in full, even if not repeated therein:

1. MIT objects to the Requests to the extent that they seek documents not relevant to the subject matter of the present lawsuit and/or are not reasonably calculated to lead to the discovery of admissible evidence.

¹ Although Harman identified its First Set of Requests as including “Nos. 1-28” under the caption and in the Certificate of Service, Harman actually served twenty-nine Requests, all of which MIT responded to or objected to. Harman’s Second Set of Requests incorrectly begins with number 27. Harman’s Second Set of Requests should be numbered 30-61. MIT’s responses herein correspond to the renumbered requests.

2. MIT objects to the Requests to the extent that they seek documents or materials that constitute or contain information protected by the attorney-client privilege, the work-product doctrine, and/or any other applicable privilege. Any disclosure of any document that is subject to such a privilege or protection is inadvertent and shall not constitute a waiver of any such privilege or protection, and any such document should, if located by Harman, be promptly returned to MIT. Any information or documents withheld on these grounds will be identified on a separate privilege log, except for such documents that were created or that came into being after initiation of this lawsuit.

3. MIT objects generally to the Requests to the extent they seek confidential documents or materials. MIT will produce confidential information only subject to the protective order entered by the Court that limits the use of documents and information produced by MIT in this action to this proceeding and protects any documents containing MIT confidential information called for by Harman's Requests from unauthorized disclosure to Harman representatives or any third parties.

4. MIT objects generally to the Requests to the extent that they attempt or purport to impose obligations on MIT beyond those required by the Federal Rules of Civil Procedure 26 and 34, or the Local Rules of Practice of the United States District Court for the District of Massachusetts.

5. MIT objects to the Requests to the extent that they seek documents already in Harman's possession, equally accessible to Harman and/or publicly available.

6. MIT objects generally to the Requests to the extent that they seek documents or materials not in MIT's possession, custody, or control.

7. MIT objects generally to the Requests to the extent that they seek to characterize evidence in this matter. To the extent that MIT expressly or impliedly adopts any term used by the Requests, such adoption is specifically limited to these responses and does not constitute an admission of fact or law in this action or any other matter.

8. None of the General Objections and Responses herein is a direct or indirect admission (i) of the truth or accuracy of any statement or characterization asserted by Harman in any pleading or other filing with the Court; (ii) of the validity of any objection or response by Harman to any discovery Request propounded in this action by MIT; or (iii) that any discovery Request propounded by MIT in this action is wholly or partially objectionable under any applicable law or rules.

9. MIT generally objects to the Requests to the extent that MIT has not completed its investigation of the facts in this case and has not completed discovery in this action. Any responses to these Requests are based on information presently known to MIT and are without prejudice to MIT's right to supplement its responses and produce as evidence, any subsequently discovered documents.

10. MIT generally objects to the Requests to the extent they are overly broad or unduly burdensome, including without limitation Requests that require production of "all" or "any" documents.

11. MIT generally objects to the Requests to the extent they are vague, ambiguous, confusing, incomprehensible and/or unanswerable because of undefined or ill-defined terms and/or confusing syntax, or they fail to describe with reasonable particularity the documents sought.

12. MIT's willingness to provide documents or things is not a concession that the subject matter of the particular document is discoverable, relevant to this action, or admissible as evidence.

13. MIT expressly reserves all objections as to relevance, authenticity, and/or admissibility of any documents produced.

14. MIT has responded to the Requests as it interprets and understands each Request made therein. If Harman subsequently asserts an interpretation of any Request that differs from the understanding of MIT, MIT reserves the right to supplement its objections and responses.

15. The presence or absence of any general or specific objection does not mean that MIT does not object on any other grounds.

16. MIT incorporates its General Objections to Harman's First Set of Requests for the Production of Documents and Things (Nos. 1-28 [renumbered 1-29]) as if fully set forth herein.

17. MIT incorporates its General Objections to Harman's First Set of Interrogatories (Nos. 1-7) and Harman's Second Set of Interrogatories (Nos. 8-20) as if fully set forth herein.

18. MIT objects to the production of any document falling within the General Objections set forth above. A statement by MIT that it will produce responsive, non-privileged, and non-work product protected documents means that MIT will produce documents, if any exist, within its possession, custody, or control that (a) are responsive to the particular Request or fall within the particular description in question, (b) have been located by MIT after a reasonable search, and (c) do not fall within one of the General Objections set forth above or within any of the objections contained in the specific responses set forth below.

19. MIT generally objects to the Requests to the extent that they are duplicative of or call for documents and things that MIT has already produced responsive to Harman's First Set of

Requests for the Production of Documents and Things (Nos. 1-28 [renumbered 1-29]) or Harman's First Set of Interrogatories (Nos. 1-7).

20. Subject to these General Objections and the Specific Objections set forth below, MIT will produce responsive documents at a time and place mutually agreed upon by the parties.

SPECIFIC OBJECTIONS AND RESPONSES

REQUEST NO. 27 [RENUMBERED 30]

All documents and things referring or relating to the work of Martin Thoone, Thoone or the Phillips CARIN system, including but not limited to, all reports or correspondence provided to MIT, James R. Davis, or Christopher Schmandt.

RESPONSE TO REQUEST NO. 27 [30]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT and are equally available to Harman. MIT further objects to this Request as duplicative of Request Nos. 12, 13, 14, 17, and 18 from Harman's First Set of Requests for the Production of Documents and Things.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 28 [RENUMBERED 31]

All documents and things referring or relating to Navigation Technologies, Inc. (NavTech), NavTeq or Karlin & Collins, Inc., or any of their systems, such as Driver Guide,

ROGUE, etc., including but not limited to all reports or correspondence provided to MIT, James R. Davis or Christopher Schmandt.

RESPONSE TO REQUEST NO. 28 [31]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT and are equally available to Harman. MIT further objects to this Request as duplicative of Request Nos. 13, 14, 17, and 18 from Harman's First Set of Requests for the Production of Documents and Things.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 29 [RENUMBERED 32]

All documents and things referring or relating to the work of Otmar Pilsak or the Bosch-Blaupunkt EVA system, including but not limited to all reports or correspondence provided to MIT, James R. Davis or Christopher Schmandt.

RESPONSE TO REQUEST NO. 29 [32]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT and are equally available to Harman. MIT further

objects to this Request as duplicative of Request Nos. 1, 6, 12, 13, 14, 17, and 18 from Harman's First Set of Requests for the Production of Documents and Things.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 30 [RENUMBERED 33]

All documents and things referring or relating to the Toyota ElectroMultivision navigation device, including but not limited to all reports or correspondence provided to MIT, James R. Davis or Christopher Schmandt.

RESPONSE TO REQUEST NO. 30 [33]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT and are equally available to Harman. MIT further objects to this Request as duplicative of Request Nos. 12, 13, 14, 17, and 18 from Harman's First Set of Requests for the Production of Documents and Things.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 31 [RENUMBERED 34]

All documents and things referring or relating to the work of Davis and Trobaugh, including but not limited to all documents and things referring to their article "Direction Assistance."

RESPONSE TO REQUEST NO. 31[34]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT and are equally available to Harman. MIT further objects to this Request as duplicative of Request Nos. 1, 6, 7, 8, and 21 from Harman's First Set of Requests for the Production of Documents and Things.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 32 [RENUMBERED 35]

All documents and things referring or relating to the public use of the Direction Assistance display in the Computer Museum in Boston or elsewhere.

RESPONSE TO REQUEST NO. 32 [35]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent it seeks documents that

are beyond the custody or control of MIT and are equally available to Harman. MIT further objects to this Request as premature to the extent it calls for a legal conclusion as to public use. MIT further objects to this Request as duplicative of Request Nos. 1, 6, 7, 8, and 21 from Harman's First Set of Requests for the Production of Documents and Things.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 33 [RENUMBERED 36]

All documents and things referring or relating to field trials or other demonstrations of the Back Seat Driver system.

RESPONSE TO REQUEST NO. 33 [36]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT and are equally available to Harman. MIT further objects to this Request as duplicative of Request Nos. 9, 10, and 11 from Harman's First Set of Requests for the Production of Documents and Things.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 34 [RENUMBERED 37]

All documents and things considered during the prosecution of the '685 patent, including, but not limited to each reference identified in the September 1990 Information Disclosure Statement.

RESPONSE TO REQUEST NO. 34 [37]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request as duplicative of Request Nos. 1, 6, 16, and 17 from Harman's First Set of Requests for the Production of Documents and Things.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 35 [RENUMBERED 38]

All documents and things referring or relating to "discourse", "discourse generation" or "discourse generators", including but not limited to presentation materials, papers, articles, correspondence, and notes.

RESPONSE TO REQUEST NO. 35 [38]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request as premature to the extent

it purports to characterize or otherwise construe the claims of the '685 patent. MIT further objects to this Request as premature to the extent it calls for a legal conclusion with respect to construction of the claims of the '685 patent. MIT further objects to this Request as not limited to a particular period of time. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity.

Subject to and without waiving the foregoing general and specific objections and in addition to responsive documents that have already been produced, MIT will provide Harman with access to publicly-available materials in MIT's library for inspection of documents responsive to this Request at an agreed-upon time. MIT further identifies as responsive to this Request the following list of materials from Mr. Schmandt's files, which MIT will make available for inspection at an agreed-upon time and place.

- Computational Models of Discourse. Brady and Berwick (eds.). (Cambridge, MA: MIT Press, 1984 (second printing));
- Readings in Natural Language Processing. Grosz, Sparck Jones, and Webber (eds.). (Morgan Kaufman, 1986);
- Pragmatics. Stephen C. Levinson. (Cambridge: Cambridge University Press, 1983);
- Text generation - using discourse strategies and focus constraints to generate natural language text. Kathleen McKeown. (Cambridge: Cambridge University Press, 1985); and
- Computational linguistics - an introduction. Ralph Grishman. (Cambridge: Cambridge University Press, 1986).

- Intentions In Communication, Cohen, Morgan & Pollack (eds.), MIT Press Cambridge, Massachusetts (1990). Pierrehumbert, Julia & Hirschberg, Julia “The Meaning of Intonational Contours In The Interpretation Of Discourse” (Chapter 14).
- Lieberman, Phillip. *Intonation, Perception & Language* (Cambridge, MA: MIT Press, 1967). Chapter 7, “Prominence, Stress, and Emphasis in American English”.
- Liberman, Mark & Prince, Alan. On Stress and Linguistic Rhythm. *Linguistic Inquiry*, 8:2 (Spring 1977), 249-336.
- Pierrehumbert, Julia & Hirschberg, Julia. The Intonational Structuring Of Discourse. *Association For Computational Linguistics*, 1985 Proceedings.
- Pierrehumbert, Julia & Hirschberg, Julia. The Intonational Structuring Of Discourse. *Association For Computational Linguistics*, 1986 Proceedings.
- Hirschberg, Julia & Litman, Diane. Now Let’s Talk About Now: Identifying Cue Phrases Intonationally. *Association For Computational Linguistics*, 1987 Proceedings.
- Grosz, Barbara & Sidner, Candace. Attention, Intentions, And The Structure Of Discourse. *Computational Linguistics*, 12:3 (July-September 1986), 175-204.
- Grosz, Barbara & Hisrschberg, Julia. Some Intonational Characteristics of Discourse Structure. *Proceedings of ICSLP-92*, I, 429-432.
- Frick, Robert. Communicating Emotion: The Role of Prosodic Features. *Psychological Bulletin*, 9:3(1985), 412-429.
- Duncan, Jr. Starkey. Some Signals And Rules For Taking Speaking Turns In Conversations. *Journal Of Personality & Social Psychology*, 23:2 (1972), 283-292.
- Cruttenden. *Intonation*. (Cambridge: Cambridge University Press, 1986). “The Functions Of Intonation” (Chapter 4).
- Clark, Herbert. *Arenas of Language Use*. (The University Of Chicago Press & Center For The Study Of Language & Information, 1992). “Definite Reference And Mutual Knowledge” (Chapter 1, with Catherine R. Marshall).
- Brown, Gillian; Currie, Karen & Kenworthy, Joanne. *Questions of Intonation* (Taylor & Francis Books, Ltd., 1980). Chapter 2, 21-39.

REQUEST NO. 36 [RENUMBERED 39]

All documents and things referring or relating to “physical connectivity”, “legal connectivity” or a “computing apparatus which distinguishes between physical and legal connectivity.”

RESPONSE TO REQUEST NO. 36 [39]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request as premature to the extent it purports to characterize or otherwise construe the claims of the '685 patent. MIT further objects to this Request as premature to the extent it calls for a legal conclusion with respect to construction of the claims of the '685 patent. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 37 [RENUMBERED 40]

All documents and things referring or relating to efforts by MIT to enforce the '685 patent.

RESPONSE TO REQUEST NO. 37 [40]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request as duplicative of Request Nos. 3, 4,

5, 12, 13, 14, and 15 from Harman's First Set of Requests for the Production of Documents and Things.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 38 [RENUMBERED 41]

All documents and things referring or relating to MIT's evaluation of a reasonable royalty as a result of any alleged infringement by Harman of the '685 patent, including but not limited to all documents and things referring or relating to reasonable royalty factors as laid out in Georgia Pacifica Corp. v. United States Plywood Corp., 318 F. Supp. 116 (S.D.N.Y. 1970).

RESPONSE TO REQUEST NO. 38 [41]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 39 [RENUMBERED 42]

All documents referring or relating to MIT's alleged proper measure of royalties for any infringement of the '685 patent, including any license agreements.

RESPONSE TO REQUEST NO. 39 [42]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 40 [RENUMBERED 43]

All documents and things upon which MIT intends to rely in support of any "damages" that MIT claims as a result of any alleged infringement by Harman of the MIT patents.

RESPONSE TO REQUEST NO. 40 [43]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 41 [RENUMBERED 44]

All documents and things that MIT intends to offer into evidence in this litigation.

RESPONSE TO REQUEST NO. 41 [44]

MIT objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request as premature as discovery is ongoing. MIT will identify its trial exhibit list when required by the Court (currently set for October 6, 2006).

REQUEST NO. 42 [RENUMBERED 45]

All documents that MIT intends to rely on within the context of any hearing or trial of this matter, including any hearing that may take place with respect to claim construction.

RESPONSE TO REQUEST NO. 42 [45]

MIT objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request as premature as discovery is ongoing. MIT will identify its documents it intends to rely on at any hearing or trial when required by the Court.

REQUEST NO. 43 [RENUMBERED 46]

All documents referred to or identified in MIT's responses to Harman's First or Second Sets of Interrogatories, all documents which those interrogatories Request to be identified, all documents relied upon or consulted in formulating MIT's responses to those interrogatories, and all documents otherwise referring or relating to MIT's responses to those interrogatories.

RESPONSE TO REQUEST NO. 43 [46]

MIT objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 44 [RENUMBERED 47]

All documents provided by MIT (or its counsel) to any expert witness(es) or consultant(s).

RESPONSE TO REQUEST NO. 44 [47]

MIT objects to this Request as vague and ambiguous to the extent that it is not limited to the present litigation. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request as premature as expert discovery is ongoing.

REQUEST NO. 45 [RENUMBERED 48]

All documents provided to MIT (or its counsel) by any expert witness(es) or consultant(s).

RESPONSE TO REQUEST NO. 45 [48]

See Response to Request 44 [47].

REQUEST NO. 46 [RENUMBERED 49]

All bills, statements or invoices provided to MIT (or its counsel) by any expert witness or consultants.

RESPONSE TO REQUEST NO. 46 [49]

See Response to Request 44 [47].

REQUEST NO. 47 [RENUMBERED 50]

All documents and things referring to or evidencing prototypes or demonstrations of the Back Seat Driver invention.

RESPONSE TO REQUEST NO. 47 [50]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request as duplicative of Request Nos. 1, 6, 7, 8, 9, 10, and 11 from Harman's First Set of Requests for the Production of Documents and Things and Interrogatories Nos. 6 & 7 from Harman's First Set of Interrogatories (Nos. 1-7).

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

EXHIBIT

26

REQUEST NO. 48 [RENUMBERED 51]

All communications between MIT, James R. Davis or Christopher Schmandt and third parties referring or relating to this litigation.

RESPONSE TO REQUEST NO. 48 [51]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request as duplicative of Request Nos. 7 and 15 from Harman's First Set of Requests for the Production of Documents and Things.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 49 [RENUMBERED 52]

All documents and things that mention, concern or discuss or are in any way related to the priority date for any claim of the '685 patent.

RESPONSE TO REQUEST NO. 49 [52]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent it seeks documents that

are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available. MIT further objects to this Request as vague or ambiguous to the extent that the Request does not specify a time period or subject matter.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 50 [RENUMBERED 53]

All documents and things that support, refer or relate to, MIT's contention that MIT's patents are valid or not invalid.

RESPONSE TO REQUEST NO. 50 [53]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request as premature to the extent that it calls for legal conclusions with respect to validity. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 51 [RENUMBERED 54]

All documents and things supporting a finding of obviousness or non-obviousness of the ‘685 patent, including any documents showing any nexus between any secondary considerations of non-obviousness and the subject matters claimed in the ‘685 patent.

RESPONSE TO REQUEST NO. 51 [54]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request as premature to the extent that it calls for legal conclusions with respect to obviousness, non-obviousness, and/or secondary considerations. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman’s possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 52 [RENUMBERED 55]

All documents and things to support any alleged diligence by, or on behalf of, Mr. Davis and/or Mr. Schmandt in reducing to practice the subject matter recited in the ‘685 patent.

RESPONSE TO REQUEST NO. 52 [55]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks

documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request as premature to the extent that it calls for legal conclusions with respect to diligence and/or reduction to practice. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 53 [RENUMBERED 56]

All documents and things that support, refer or relate to MIT's contention that the '685 patent is enforceable or not unenforceable.

RESPONSE TO REQUEST NO. 53 [56]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request as premature to the extent that it calls for legal conclusions with respect to enforceability of the '685 patent. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 54 [RENUMBERED 57]

All documents and things that support, or refer or relate to, MIT's contention that Harman's products and/or process allegedly infringe the '685 patent.

RESPONSE TO REQUEST NO. 54 [57]

MIT objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request as premature to the extent that it calls for legal conclusions with respect to infringement. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 55 [RENUMBERED 58]

All documents and things that support, or refer or relate to, MIT's contention that Harman willfully infringed the '685 patent.

RESPONSE TO REQUEST NO. 55 [58]

MIT objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request as premature to the extent that it calls for legal conclusions with respect to willful infringement. MIT further objects to this Request to the extent it seeks documents that are beyond the custody or control of MIT. MIT further objects to this Request to the extent that any

responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 56 [RENUMBERED 59]

All documents and things that support, or refer or relate to, MIT's contention that MIT is entitled to damages for Harman's alleged infringement of the '685 patent.

RESPONSE TO REQUEST NO. 56 [59]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 57 [RENUMBERED 60]

All documents and things having any tendency to prove or disprove each and every affirmative defense asserted by MIT in its responsive pleadings to Harman's counterclaims.

RESPONSE TO REQUEST NO. 57 [60]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity. MIT further objects to this Request to the extent that any responsive documents are already in Harman's possession, readily accessible to Harman and/or publicly available.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

REQUEST NO. 58 [RENUMBERED 61]

All documents and things having any tendency to prove or disprove each and every allegation not admitted by MIT in its responsive pleadings to Harman's Counterclaims.

RESPONSE TO REQUEST NO. 58 [61]

MIT objects to this Request as overbroad, unduly burdensome, and unlikely to lead to the discovery of admissible evidence. MIT further objects to this Request to the extent it seeks documents protected by attorney-client privilege, the work-product doctrine or any other privilege or immunity.

Subject to and without waiving the foregoing general and specific objections, MIT states that it has already produced or logged on its privilege log any documents in its possession, custody, or control responsive to this Request.

Dated: April 21, 2006

Respectfully submitted,

Massachusetts Institute of Technology,

By its Attorneys,

/s/ John W. Pint

Steven M. Bauer (BBO# 542531)
Kimberly A. Mottley (BBO# 651190)
John W. Pint (BBO# 660548)
PROSKAUER ROSE LLP
One International Place
Boston, Massachusetts 02110-2600
Phone: 617-526-9600
Fax: 617-526-9899

CERTIFICATE OF SERVICE

I HEREBY CERTIFY that on April 21, 2006, I caused a true and correct copy of MIT'S RESPONSES TO HARMAN'S SECOND SET OF REQUESTS FOR THE PRODUCTION OF DOCUMENTS AND THINGS (NOS. 30-61) to be served on the following counsel of record via email:

Robert J. Muldoon, Jr.
SHERIN AND LODGEN, LLP
101 Federal Street
Boston, MA 02110

William A. Streff Jr., P.C.
Michelle A. H. Francis
Craig D. Leavell
Ann H. Chen
Colleen M. Garlington
KIRKLAND & ELLIS LLP
200 East Randolph Drive
Chicago, IL 60601
(312) 861-2000 (phone)
(312) 861-2200 (fax)

By: /s/ John W. Pint
John W. Pint

EXHIBIT

27

KIRKLAND & ELLIS LLP
AND AFFILIATED PARTNERSHIPS

200 East Randolph Drive
Chicago, Illinois 60601

William A. Streff, Jr., P.C.
To Call Writer Directly:
312.861.2126
wstreff@kirkland.com

312.861.2000
www.kirkland.com

Facsimile:
312.861.2200

October 10, 2006

By Electronic Mail

Steven M. Bauer, Esquire
Proskauer Rose LLP
One International Place
Boston, MA 02110-2600

Re: *MIT v. Harman*, Case No. 05-10990 DPW (D. Mass.)

Dear Steve:

Pursuant to Fed. R. Civ. P. 26(e), we are producing the enclosed documents bearing Bates numbers HAR 709859-709907. We recently received these documents from Mr. Gregory Grove, who a short time ago joined our firm as a partner in our corporate department. Mr. Grove was an undergraduate student at MIT in the 1988-92 time period. Upon reviewing the attached documents, we believe that they strongly support Harman's positions that the '685 patent is (1) invalid under 35 U.S.C. §102(b) and (2) unenforceable due to inequitable conduct before the Patent Office. Upon reviewing MIT's discovery responses in view of the existence of these documents and the events they depict, we believe that we have further support for our anticipated dispositive motion for summary judgment and for our petition for attorneys' fees.

Specifically, attached is an April 10, 1989 Undergraduate Research Opportunity Program Final Paper ("UROP") authored by Mr. Grove, under the direction of MIT's Christopher Schmandt (a named inventor of the patent-in-suit), which states *inter alia* that "[t]he Backseat Driver is a system for maintaining a moving vehicle's position, speed, and direction, and directing the operator of the vehicle to a specific destination." See HAR 709860. Another of these documents shows that "Jim Davis" (a named inventor of the patent-in-suit) presented his "Thesis Defense" — "Telling You Where to Go" via "Back Seat Driver" — on May 26, 1989. See HAR 709861 ("The Back Seat Driver is a computer program which uses synthetic speech to give instructions to the driver of the car as needed while driving").

Steven M. Bauer, Esq.

October 10, 2006

Page 2

These documents, separately and together, are inconsistent with MIT's position that the earliest reduction to practice of any claims of the '685 patent was June 1989. *See* MIT's June 16, 2006 Supplemental Response to Harman's Interrogatory No. 8.¹ They also are inconsistent with MIT's position that the Back Seat Driver prototypes identified in the '685 specification were "created for experimentation, research and development purposes." *See* MIT's Amended Response to Harman's Interrogatory No. 6.² Instead, these documents show that Mr. Davis' thesis was complete and the inventions therein were ready for patenting no later than May 1989. *See Pfaff v. Wells Electronics, Inc.*, 525 U.S. 55, 57, 119 S. Ct. 304, 311-12 (1998) (holding that the public use bar of § 102(b) invalidates a patent where an invention was publicly disclosed more than one year before the patent application date, and was ready for patenting either by reduction to practice or because the "inventor had prepared drawings or other descriptions of the invention that were sufficiently specific to enable a person skilled in the art to practice the invention"). Under *Pfaff*, Jim Davis' public defense of his thesis was a sufficiently specific public disclosure to invalidate all of the claims in the '685 patent, particularly given that MIT incorporated Davis' thesis by reference into the patent and Davis' thesis largely comprises the patent's specification. *Id.*

Moreover, MIT's failure to disclose these documents or their substance during prosecution of the '685 patent calls into question MIT's candor with the U.S. Patent Office. When the examiner initially rejected all claims "under 35 U.S.C. § 102(e) as being clearly anticipated by the Ph.D. thesis of James Raymond Davis," *see* File Wrapper ("FW") at 97, because "the title page of the thesis bears a submission date of August 4, 1989, more than one year before the filing date of the present application," *see* F.W. at 110, MIT swore:

- "August 4 [1989] is the date the thesis was signed, and not the date on which the thesis became available to the public."
- "the thesis ... is therefore not 102 art with respect to the present application."

¹ MIT states that the following claims were reduced to practice at least as early as June of 1989: 1-4, 7-9, 21, 23-25, 32-37, 40, 41, 53, 55, 57, and 58. MIT further states that the following claims were reduced to practice at least as early as August 4, 1989: 10-12, 15, 19, 20, 26-31, 42-49, 51, 52, and 54. MIT further states that the following claims were reduced to practice at least as early as the filing date of the '685 patent, August 9, 1990: 5, 6, 13, 14, 16-18, 20, 22, 38, 39, 50, and 56."

² MIT wrote: "All three test units were created for experimentation, research and development purposes in conjunction with research sponsored by NEC. The persons with information regarding the development of the device and any associated disclosure are the inventors, Dr. Davis and Prof. Schmandt."

Steven M. Bauer, Esq.

October 10, 2006

Page 3

F.W. at 110. In direct violation of its duty of candor owed to the Patent Office, MIT misrepresented and withheld critical facts concerning MIT's first public disclosure of the Back Seat Driver and, thus, the inventions claimed in the '685 patent.

There can be no reasonable justification for withholding this information during prosecution or during this litigation. These documents are responsive to several of Harman's document requests and interrogatories, as well as some deposition questions.³ In light of MIT's "eye toward litigation" extending back to 1989, MIT has had a continuing obligation to preserve these documents. *See July 7, 2006 MIT's Opposition To Harman's Motion To Compel Discovery Of Documents MIT Has Identified As Protected On Its Privilege Logs at 6; see also* MIT's July 7, 2006 Privilege Log Entry Nos. 37, 73, 76, 76, 96 (all dated in 1989). MIT should have produced them (or any documents disclosing this information, which MIT did not do) to Harman at the very outset of this case. MIT also should have identified in its initial disclosures the individuals who performed an UROP as part of "back seat driver" since each undoubtedly has knowledge concerning some or all of the inventions claimed in the '685 patent. MIT's inexcusable failure to do any of this, at a minimum, lays a strong foundation for Harman's dispositive motion and fee petition.

However, prior to filing its motion, Harman is willing to renew the settlement offer it presented prior to the August 4, 2005 Scheduling Conference. Specifically, Harman will stipulate to a dismissal with prejudice of MIT's complaint and a dismissal without prejudice of Harman's counter-claims, provided MIT fully and forever releases and discharges any and all claims against Harman, with no royalties due and with MIT to bear its own and Harman's legal fees incurred to date in defending against MIT's complaint. We trust MIT will consider Harman's offer and respond to us in the near future.

³ See, e.g., Harman's Document Request No. 1 (seeking all "documents and things referring or relating to U.S. Patent No. 5,177,685 ... the application that lead to its issuance ... or the subject matter thereof, including without limitation all engineering notebooks, laboratory notebooks, records, logs, and files created or maintained by, or at the direction of, Dr. James R. Davis, Mr. Christopher M. Schmandt, or Dr. Sam Pasternack"); Request No. 27 (seeking all "thesis drafts, revisions, notes and documents referring or relating or leading up to the September 1989 thesis paper..."); Request No. 6 (seeking all documents "referring or relating to Back Seat Driver"); Request No. 7 (seeking all documents referring or relating to "any disclosure of any portion of the subject matter" of the patent); Request No. 8 (seeking all documents referring or relating to "any disclosure to, or use by any third party of any portion of the subject matter of the [patent] before the filing date..."); Request No. 9 (seeking all documents regarding the "first, use, sale or display of any alleged invention" of the patent); Request No. 10 (seeking all documents regarding "any experimental use, field test, and/or prototype"); Request No. 17 (seeking all documents regarding any prior art "known, considered, reviewed or evaluated"); and Request No. 22 (seeking the "personnel, student and/or employment history files for each of the named inventors ... and any MIT students and/or employees" involved).

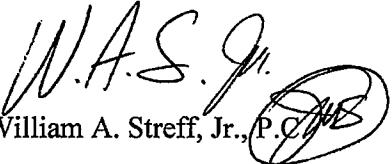
Steven M. Bauer, Esq.

October 10, 2006

Page 4

Please contact me with any questions, or if you would like to discuss this matter further.

Very truly yours,


William A. Streff, Jr., P.C.


cc: Robert J. Muldoon, Esq. (via E-Mail)

EXHIBIT

28

UNITED STATES DISTRICT COURT
DISTRICT OF MASSACHUSETTS

MASSACHUSETTS INSTITUTE)
OF TECHNOLOGY,)
)
Plaintiff,) CIVIL ACTION
) NO. 05-10990-DPW
v.)
)
HARMAN INTERNATIONAL)
INDUSTRIES, INC.,)
)
Defendant.)

**ORDER REGARDING REVIEW OF DOCUMENTS
SUBMITTED FOR IN CAMERA INSPECTION**

This court has reviewed “MIT’s (Second) In Camera Submission of Documents Related to Harman and NEC Over Which MIT Continues to Assert Work Product and Attorney-Client Privilege Protections” as well as “Harman’s August 21, 2006 Submission for In Camera Review Pursuant to the Court’s August 4, 2006 Order.” In addition, this court has reviewed document number 471 listed on MIT’s privilege log, which MIT had submitted to the court on August 8, 2006. Based on this review, this court concludes as follows:

1. The following documents identified on MIT’s privilege log were properly withheld from production pursuant to the work product doctrine: Nos. 329, 337, 471, 558, 559, 562, 564, 566, 567, 568, 605, 762 and 782. MIT has shown that these documents were prepared by a party and “can be fairly said to have been prepared or

obtained because of the prospect of litigation” or in connection with ongoing litigation.

In re Grand Jury Subpoena, 220 F.R.D. 130, 146 (D. Mass. 2004).

2. The following documents identified on MIT’s privilege log were properly withheld from production based on the attorney-client privilege: Nos. 280 and 336 (redacted portions). MIT has demonstrated that these documents consist of or describe confidential communications between MIT and its counsel “for the purpose of seeking, obtaining, or providing legal assistance to the client.” E. Epstein, The Attorney-Client Privilege and the Work Product Doctrine (4th ed.) at 46 (ABA 2001). See also City of Worcester v. HCA Mgmt. Co., Inc., 839 F. Supp. 86, 88 (D. Mass. 1993) (describing privilege).

3. MIT has not established that the community of interest or common interest doctrine provides an exception to the waiver of the attorney-client privilege with respect to documents that were disclosed to NEC. As an initial matter, nothing in the NEC related documents or in MIT’s submissions demonstrates that MIT and NEC were joint clients of MIT’s patent counsel during the prosecution of the ‘685 patent or at any other time. See Massachusetts Eye & Ear Infirmary v. QLT Phototherapeutics, Inc., 412 F.3d 215, 228 (1st Cir. 2005) (in evaluating common interest doctrine, proper focus is on the joint representation of the parties by one party’s attorney), cert. denied, 126 S. Ct. 2292, 164 L. Ed. 2d 814 (2006); Mass. Eye & Ear Infirmary v. QLT Phototherapeutics, Inc., 167 F. Supp. 2d 108, 117-124 (D. Mass. 2001) (analyzing objective reasonableness of party’s belief that joint attorney-client relationship existed for purposes of applying

common interest doctrine). See also In re Regents of the Univ. of Cal., 101 F.3d 1386, 1389-90 (Fed. Cir. 1996) (evaluating, for purposes of applying the attorney-client privilege, whether the same attorney was representing two entities on the same matter and whether the clients shared a common legal interest). Rather, the documents indicate that MIT's counsel at first represented only MIT, without any involvement of NEC, and later communicated with NEC's own counsel for informational purposes only. Thus, the record does not support the conclusion that there was a professional relationship between MIT's counsel and NEC, or that NEC consulted with MIT's counsel for the purpose of seeking legal advice. Id. at 1390. Thus, MIT has not "met its burden of establishing that [NEC] shared an attorney-client relationship with [MIT's counsel] on the preparation and prosecution of the [relevant patent] application." Mass Eye & Ear Infirmary, 167 F. Supp. 2d at 117-18.

Moreover, MIT has not established that it shared a common legal interest with NEC. In particular, with the exception of document no. 7, all of the documents at issue appear to post-date the license agreement between MIT and NEC.¹ Therefore, although NEC was a sponsor of the technology that MIT was seeking to patent and had an option to obtain an exclusive license, at the time when the documents were generated NEC had only a non-exclusive, royalty-free non-transferable license. Accordingly, MIT has not shown that NEC shared its interest in obtaining a valid and enforceable patent. See

¹ Although it is unclear when NEC entered into the license agreement, MIT executed the agreement on May 28, 1991.

Regents, 101 F.3d at 1390 (in contrast to a situation where the parties' legal interests were "substantially identical because of the potentially and ultimately exclusive nature of the [parties'] license agreement," a patentee and a nonexclusive licensee do not share identical interests). Consequently, the following documents identified on MIT's privilege log must be produced in their entirety: Nos. 3, 40, 41, 68, 84, 116, 120, 122 and 128. Additionally, MIT must produce the following documents in redacted form in order to protect the attorney-client communications contained therein: Nos. 6, 7 and 663. Document No. 212 was properly withheld based on the attorney-client privilege. Finally, document No. 17 was not disclosed to NEC and need not be disclosed unless it is responsive.

4. Harman has demonstrated that all of the documents that it submitted to this court for in camera review on August 21, 2006 were properly withheld from production pursuant to the attorney-client privilege.

/ s / Judith Gail Dein

Judith Gail Dein
United States Magistrate Judge

DATED: September 27, 2006

EXHIBIT

29

Westlaw.

--- F.3d ----

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

(Cite as: --- F.3d ----)

Page 1

Briefs and Other Related Documents

eSpeed, Inc. v. BrokerTec USA, L.L.C.C.A.Fed. (Del.), 2007. Only the Westlaw citation is currently available.

United States Court of Appeals, Federal Circuit.
eSPEED, INC., Cantor Fitzgerald, L.P., CFPN,
L.L.C., and eSpeed Government Securities, Inc.,

Plaintiffs-Appellants,

v.

BROKERTEC USA, L.L.C., Garban, L.L.C., OM
Technology US, Inc., and OM Technology AB, De-
fendants-Appellees.

No. 2006-1385.

March 20, 2007.

Background: Owner of patent for electronic trading platform used in trading government securities sued competitor for infringement. Competitor counter-claimed that patent was unenforceable, and that suit was barred by laches. The United States District Court for the District of Delaware, Kent A. Jordan, J., 417 F.Supp.2d 580, found patent was unenforceable because of inequitable conduct, and entered final judgment in favor of competitor. Patent owner appealed.

Holdings: The Court of Appeals, Moore, Circuit Judge, held that:

(1) false statements in declarations submitted during prosecution, to disclose prior art and attempt to cure the earlier non-disclosure of this prior art, were material, and

(2) finding that declarations were submitted with intent to deceive the Patent and Trademark Office (PTO) was not clearly erroneous.

Affirmed.

[1] Patents 291 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings

Therein in General. Most Cited Cases

Inequitable conduct while prosecuting patent applications in the Patent and Trademark Office (PTO) includes affirmative misrepresentation of a material fact, failure to disclose material information, or submission of false material information, coupled with an intent to deceive.

[2] Patents 291 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings

Therein in General. Most Cited Cases

In cases where the omission or misrepresentation while prosecuting patent applications in the Patent and Trademark Office (PTO) is highly material, less evidence of intent will be required to find that inequitable conduct has occurred.

[3] Patents 291 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings

Therein in General. Most Cited Cases

Once a district court has found a threshold level of both materiality and intent to deceive, the district court must balance the evidence to determine if equity should render the patent unenforceable.

[4] Patents 291 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings

Therein in General. Most Cited Cases

Patents 291 324.54

291 Patents

291XII Infringement

291XII(C) Suits in Equity

291k324 Appeal

291k324.54 k. Presumptions and Discre-

tion of Lower Court. Most Cited Cases

The ultimate conclusion that a patent is unenforce-

--- F.3d ----

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

(Cite as: --- F.3d ----)

Page 2

able is an equitable decision committed to the discretion of the district court, that the Court of Appeals reviews for an abuse of discretion.

[5] Federal Courts 170B 812

170B Federal Courts

170BVIII Courts of Appeals

170BVIII(K) Scope, Standards, and Extent

170BVIII(K)4 Discretion of Lower Court

170Bk812 k. Abuse of Discretion. Most Cited Cases

To overturn a discretionary ruling of a district court, an appellant must establish: (1) that the ruling is based upon clearly erroneous findings of fact; (2) that the ruling is a misapplication or misinterpretation of applicable law; or (3) that the ruling evidences a clear error of judgment on the part of the district court.

[6] Patents 291 324.55(2)

291 Patents

291XII Infringement

291XII(C) Suits in Equity

291k324 Appeal

291k324.55 Questions of Fact, Verdicts, and Findings

291k324.55(2) k. Clearly Erroneous Findings. Most Cited Cases

Court of Appeals reviews the underlying facts of materiality of an omission or misrepresentation while prosecuting patent applications in the Patent and Trademark Office (PTO) and intent for clear error; under this standard, the Court will not disturb factual findings unless it has a definite and firm conviction that a mistake has been committed.

[7] Patents 291 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings Therein in General. Most Cited Cases

False statements in declarations submitted during prosecution of patent for electronic trading platform used in trading government securities, to disclose prior art and attempt to cure the earlier non-disclosure of this prior art, were material; declaration, stating prior art did not include "new rules" of workup for second-

ary market trades for fixed income securities left examiner with the impression that examiner did not need to conduct any further investigation, and this information would have been important to a reasonable examiner in deciding whether to allow application to issue as a patent.

[8] Patents 291 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings Therein in General. Most Cited Cases

Under the reasonable examiner standard, information is "material" when a reasonable examiner would consider it important in deciding whether to allow the application to issue as a patent.

[9] Patents 291 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings Therein in General. Most Cited Cases

False statements are more likely material when embodied in declarations or affidavits submitted to the Patent and Trademark Office (PTO).

[10] Patents 291 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings Therein in General. Most Cited Cases

To satisfy the intent to deceive element of inequitable conduct while prosecuting patent applications in the Patent and Trademark Office (PTO), the involved conduct, viewed in light of all the evidence, including evidence of good faith, must indicate sufficient culpability to require a finding of intent to deceive.

[11] Patents 291 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings Therein in General. Most Cited Cases

There is no requirement that intent to deceive while prosecuting patent applications in the Patent and

--- F.3d ----

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

(Cite as: --- F.3d ----)

Page 3

Trademark Office (PTO) be proven by direct evidence; intent to deceive may be inferred from the facts and circumstances surrounding the applicant's overall conduct.

[12] Patents 291 ↪ 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings Therein in General. Most Cited Cases
An inference of intent to deceive may arise where material false statements are proffered in a declaration or other sworn statement submitted to the Patent and Trademark Office (PTO).

[13] Patents 291 ↪ 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings Therein in General. Most Cited Cases
Finding that applicants' declarations as to application for patent for electronic trading platform used in trading government securities were submitted with intent to deceive the Patent and Trademark Office (PTO) was not clearly erroneous, where applicants submitted declarations in an apparent attempt to purge possible inequitable conduct in first application and to disclose prior art in continuing application, and instead of being candid, declarations disingenuously stated that prior art did not include new rules of workup for secondary market trades for fixed income securities.

[14] Patents 291 ↪ 97

291 Patents

291IV Applications and Proceedings Thereon

291k97 k. Patent Office and Proceedings Therein in General. Most Cited Cases
The affirmative act of submitting an affidavit must be construed as being intended to be relied upon, for purposes of satisfying the intent to deceive element of inequitable conduct while prosecuting patent applications in the Patent and Trademark Office (PTO).

Patents 291 ↪ 328(2)

291 Patents

291XIII Decisions on the Validity, Construction, and Infringement of Particular Patents

291k328 Patents Enumerated

291k328(2) k. Original Utility. Most Cited

Cases

5,905,974, 6,560,580. Cited.

Gary A. Rosen, Law Offices of Gary A. Rosen, P.C., of Philadelphia, PA, argued for plaintiffs-appellants. With him on the brief was Jack B. Blumenfeld, Morris, Nichols, Arsh & Tunnell LLP, of Wilmington, DE.

J. Michael Jakes, Finnegan, Henderson, Farabow, Garrett & Dunner, L.L.P., of Washington, DC, argued for defendants-appellees. With him on the brief was James R. Barney.

Before LINN, DYK, and MOORE, Circuit Judges.
MOORE, Circuit Judge.

*1 Plaintiffs-appellants eSpeed, Inc., Cantor Fitzgerald, L.P., CFPH, L.L.C. and eSpeed Government Securities, Inc. (collectively Cantor) appeal from the district court's final judgment declaring claims 20-23 of United States Patent No. 6,560,580 (the 580 patent) invalid, declaring the 580 patent unenforceable, and entering final judgment in favor of defendants. Because we affirm the district court's conclusion that the 580 patent is unenforceable due to inequitable conduct, we need not decide the other issues raised by Cantor.

I. BACKGROUND

A. Technological Background

1. Methods involving "open outcry" and "trade capture"

The 580 patent pertains to automated methods and systems for trading financial instruments, particularly fixed income securities. Prior to the development of the invention described in the 580 patent, financial instruments were sold using an "open-outcry" method whereby "voice brokers" would express various bid and offer prices for a given instrument. 580 patent, col.3 ll.6-9. According to the 580 patent, "[t]his expression would involve the loud oral 'cry' of a cus-

--- F.3d ----

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

(Cite as: --- F.3d ----)

Page 4

tomer-proposed bid or offer and the coordination with the fellow representatives regarding the extraction of complementary positions-until a transaction match is made and a deal is done." *Id.* at col.3 ll.9-12. Open outcry auction bond brokering served its customers well because it was efficient and permitted trading at "near perfect market pricing." *Id.* at col.2 ll.64-66.

While voice brokers were participating in open outcry trading, a process known as "trade capture" was performed by designated clerks. *Id.* at col.3 ll.13-18. These clerks would attempt to record the "outcry of many individual brokers simultaneously" using electronic input devices, such as a computer or workstation. *Id.* at col.3 ll.18-21. As might be apparent from its description, the quality of the information inputted into the electronic devices by a clerk was "a function of the interpretative skill of the input clerk, and the volume and the volatility of customer orders." *Id.* at col.3 ll.22-23.

The inventors of the 580 patent recognized that there was a need for greater efficiency and accuracy in the trading of instruments such as fixed income securities. *Id.* at col.3 l.29-col.4 l.3. Therefore, they sought to create a system to automate the trading process and avoid the use of open outcry and trade capture processes.

2. "New" vs. "Old" Rules of Workup

Traders in the secondary market for fixed income securities, such as United States Treasuries (e.g., T-bills, notes, and bonds), do not want to reveal the full volume that they are willing to trade at a given price because this information might be used against the trader by other market participants. However, in order to foster liquidity in the market, customers who initiate the trade at a given price are provided with the exclusive option to incrementally increase their purchase volume. This exclusivity is known as "workup rights." When given exclusivity, a customer can gradually increase the volume of his purchase while determining how the market is reacting to the purchase before trading further.

*2 In what the parties have referred to as the "old

rules" of workup, when the first buyer or seller has completed his transaction, new buyers or sellers are sequentially given, in the order in which they expressed interest, exclusive workup rights. One problem with the old rules of workup was that a few participants could tie up the market for long periods of time. As a result, brokers would, on occasion, engage in side deals to avoid losing business. Cantor presented evidence at trial showing that the old rules led to "chaos" and "pandemonium" when trading volume was heavy.

Because of the problems associated with the old rules of workup, Cantor employees began to develop new rules of workup. In 1994, the new rules of workup were designed to provide exclusivity to an initial pair of market participants, in a manner similar to the old rules of workup. After the initial pair of traders was finished, orders that were placed while the initial pair had exclusivity would then be rapidly executed in time priority order. Thus, by limiting exclusivity to the first pair of traders, the new rules of workup still provided an incentive for the first pair of traders to create liquidity while at the same time avoiding a long queue of traders waiting for their chance to trade.

3. The Super System/CFTS2.0

In the late 1980's Cantor began looking to replace its decade-old trade capture system with a new system. Between 1987 and 1992, programmers and software developers at Cantor wrote software code that would later become known internally as the "Super System" or the Cantor Fitzgerald Trading System (CFTS) 2.0. The district court found that "the Super System would provide a platform to support both automated trading and traditional outcry trading [using trade capture]." *eSpeed, Inc. v. BrokerTec USA, L.L.C. (Unenforceability Ruling)*, 417 F.Supp.2d 580, 586 (D.Del.2006). The Super System included software code for various trading states including a workup state.^{FN1} The Super System also included code that allowed brokers to use either the old rules or the new rules.^{FN2}

As early as 1993, the Super System was used in Cantor trading rooms to conduct trades. *Id.* at 588. After

--- F.3d ---

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

(Cite as: --- F.3d ----)

Page 5

using the Super System in 1993, Cantor determined that the system was too slow to be commercially used as an automated trading system and used it solely “as an order entry system, to support open outcry trading.” *Id.* The Super System was used as a trade capture system in Cantor’s trading rooms between 1993 and 1995 to transact billions of dollars worth of trades.

4. Subsequent Versions of CFTS

Cantor continued to improve the Super System in an attempt to create a more efficient automated trading system and first used Super System’s successor, CFTS 3.1, in Cantor’s Long Bond Room in 1995 during the week between Christmas and New Year’s Day. *Unenforceability Ruling*, 417 F.Supp.2d at 588. CFTS 3.1 was able to implement trades efficiently enough that Cantor deemed it to be commercially viable.

B. The Prosecution History

*3 The 580 patent was filed as United States Patent Application No. 09/294,526 (the 526 application). The 526 application claimed priority under 35 U.S.C. § 120 to United States Patent Application No. 08/766,733 (the 733 application). The 733 application was filed on December 13, 1996. This patent application matured into United States Patent No. 5,905,974 (the 974 patent) on May 18, 1999. Neither the Super System nor any other version of CFTS was disclosed to the United States Patent and Trademark Office (PTO) during the course of prosecution of the 733 parent application. *Unenforceability Ruling*, 417 F.Supp.2d at 588.

Shortly after the 974 patent issued, Cantor asserted the 974 patent against Liberty Brokerage. In preparing for that lawsuit, Cantor’s outside counsel learned about the Super System and realized that it had not been disclosed to the PTO. Cantor dismissed its suit against Liberty Brokerage after this discovery.

In an effort to purge the possible inequitable conduct with regard to the 974 patent and avoid a similar problem with any patent that might issue based on the 526 application, Cantor submitted three declarations and numerous exhibits purporting to describe the Su-

per System to the patent examiner in connection with the 526 application. Declarations were submitted by Stuart A. Fraser, Vice Chairman of Cantor Fitzgerald Securities (CFS), Howard W. Lutnick, Chairman of CFS, and Bijoy Paul, a Cantor employee responsible for the development of Super System’s successors. Each of these individuals averred that he was an inventor of the 526 application and that they did not realize that they were under a duty to disclose the Super System during the prosecution of the 733 application. One declaration, submitted by Paul, stated that “[t]he Super System … *did not include new rules*” and that “[t]he Super System code was based on ‘old rules’ in which each successive broker had a period of exclusive control over the trade.” Decl. of Bijoy Paul, 526 application, ¶¶ 11, 20 (Jan. 31, 2002) (emphasis added). The declarations submitted by Cantor referenced various exhibits that were submitted to the PTO along with the declarations. The exhibits amassed to 1139 pages, and included portions of the Super System source code.

The patent examiner considered the materials submitted by Cantor and concluded that the cited papers (presumably referring to the exhibits which were internal Cantor documents) did not constitute prior art, but were “given due consideration.”

C. The District Court Proceedings

Cantor filed suit against BrokerTec USA, L.L.C., Garban, L.L.C., OM Technology US, Inc. and OM Technology AB (collectively BrokerTec) in the United States District Court for the District of Delaware on June 30, 2003. Cantor asserted that BrokerTec infringed claims 20-23 of the 580 patent. During the course of the litigation, the district court ruled, on motion by BrokerTec, that the submission of the inventor’s declarations during the prosecution of the 526 application waived the attorney-client privilege with respect to discussions between Cantor’s attorneys and the inventors regarding the Super System. Accordingly, the district court ordered the production of certain documents related to those discussions. See, e.g., *eSpeed, Inc. v. BrokerTec USA, L.L.C.*, No. 03-612-KAJ, 2004 WL 1812702 (D.Del. Aug. 5, 2004).

--- F.3d ----

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

(Cite as: --- F.3d ----)

Page 6

*4 The district court construed the claims on September 9, 2004, *eSpeed, Inc. v. BrokerTec USA, L.L.C.*, No. 03-612-KAJ, 2004 WL 2346141, at *8 (D.Del. Sept.9, 2004), and granted BrokerTec's motion for partial summary judgment of non-infringement under the doctrine of equivalents, concluding that an amendment to the claims created prosecution history estoppel with respect to the alleged equivalents. *See eSpeed, Inc. v. BrokerTec USA, L.L.C.*, 342 F.Supp.2d 244, 250-51 (D.Del.2004).

A jury trial commenced on February 7, 2005, with inequitable conduct tried simultaneously to the district court. The jury returned a verdict of infringement in favor of Cantor, but held that the patent was invalid due to lack of written description under 35 U.S.C. § 112, ¶ 2.

Shortly after denying Cantor's JMOL motion, the district court held that the 580 patent was unenforceable because of inequitable conduct. *See Unenforceability Ruling*, 417 F.Supp.2d 580. The district court determined that there were two separate grounds for unenforceability and that either, standing alone, rendered the 580 patent unenforceable. *Id.* at 599. With respect to the first ground of inequitable conduct, the district court determined that the use of the Super System more than one year prior to the filing date of the 733 application was material prior art that should have been disclosed during prosecution of the 733 application. *Id.* at 589-90, 593. The district court concluded that the inventors intended to deceive the PTO because each of the inventors had a significant role in creating the Super System and because of the Super System's high materiality. *Id.* at 594. Moreover, the district court found further evidence of intent when, prior to filing the 733 application, Lutnick stated that he had "wanted to Patent [Cantor's] Super System & its rules (in general) for over a year." *Id.* The district court found that inequitable conduct during the prosecution of the 974 patent infected the prosecution of the 580 patent because the patents were closely related and Cantor failed to cure the inequitable conduct by filing the inventor declarations in the 526 application. *See id.* at 595-96.

As the second basis for rendering the 580 patent unenforceable, the district court concluded that the three

inventor declarations submitted to the PTO to disclose the Super System included material misrepresentations. *Id.* at 597. Specifically, the district court concluded that the declarations included statements "that the Super System did not contain any code for the 'new rules' of trading." *Unenforceability Ruling*, 417 F.Supp.2d at 597. This, the district court found, was not true. *Id.* The district court concluded that although there might have been different "new rules" at different times, all "new rules" had something in common: they limited exclusivity. *Id.* This was "perhaps the primary object" of the invention claimed in the 580 patent. *Id.* Thus, the district court found the misrepresentations in the declarations to be material. *Id.* at 598. The district court found sufficient facts to infer an intent to deceive, based in part on the fact that the declarations were worded in such a way to make the examiner believe that there were no "new rules" in the Super System. *Id.* Weighing the materiality with intent to deceive, the district court held the 580 patent unenforceable. *Id.*

D. The Present Appeal

*5 On appeal, Cantor challenges the district court's claim construction, the district court's decision regarding the applicability of the doctrine of equivalents, the district court's denial of JMOL on the issue of lack of written description, the district court's determination that the 580 patent is unenforceable, and that the district court's decisions regarding the existence and scope of waiver of the attorney-client privilege based on the submission of the inventor's declarations warrants a new trial. We have jurisdiction under 28 U.S.C. § 1295(a)(1).

We conclude that the district court did not clearly err in finding that the declarations included material false statements and were submitted with an intent to deceive. Because the district court did not abuse its discretion in rendering the 580 patent unenforceable, we need not reach the remaining issues raised by Cantor.

II. DISCUSSION

[1][2][3] "[I]nequitable conduct includes affirmative misrepresentation of a material fact, failure to disclose material information, or submission of false

--- F.3d ----

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

(Cite as: --- F.3d ----)

Page 7

material information, coupled with an intent to deceive.” *Pharmacia Corp. v. Par Pharm., Inc.*, 417 F.3d 1369, 1373 (Fed.Cir.2005) (quoting *Molins PLC v. Textron, Inc.*, 48 F.3d 1172, 1178 (Fed.Cir.1995)). In cases where the omission or misrepresentation is highly material, “less evidence of intent will be required in order to find that inequitable conduct has occurred.” *PerSeptive Biosystems, Inc. v. Pharmacia Biotech*, 225 F.3d 1315, 1319 (Fed.Cir.2000). Once a district court has found a threshold level of both materiality and intent to deceive, the district court must balance the evidence to determine if equity should render the patent unenforceable. *See LaBounty Mfg., Inc. v. U.S. Int'l Trade Comm'n*, 958 F.2d 1066, 1070 (Fed.Cir.1992).

[4][5] The ultimate conclusion that a patent is unenforceable is an equitable decision committed to the discretion of the district court that we review for an abuse of discretion. *See Flex-Rest, L.L.C. v. Steelcase, Inc.*, 455 F.3d 1351, 1357 (Fed.Cir.2006). “To overturn a discretionary ruling of a district court, the appellant must establish that the ruling is [1] based upon clearly erroneous findings of fact or [2] a misapplication or misinterpretation of applicable law or that [3] the ruling evidences a clear error of judgment on the part of the district court.” *Kingsdown Med. Consultants, Ltd. v. Hollister, Inc.*, 863 F.2d 867, 876 (Fed.Cir.1988) (en banc in relevant part).

[6] We review the underlying facts of materiality and intent for clear error. *Agfa Corp. v. Creo Prods., Inc.*, 451 F.3d 1366, 1371 (Fed.Cir.2006). Under this standard, we will not disturb factual findings “unless we have a definite and firm conviction that a mistake has been committed.” *Flex-Rest*, 455 F.3d at 1357 (quoting *Critikon v. Becton Dickinson Vascular Access, Inc.*, 120 F.3d 1253, 1255 (Fed.Cir.1997)).

A. Materiality

[7] Cantor contends that the declarations submitted during the prosecution of the 580 patent were not material. Cantor first argues that the inventors’ statements are immaterial because the record does not reflect that any brokers ever used the code for the new rules of trading when the Super System was deployed in Cantor’s trading rooms. Second, Cantor argues that

the exhibits submitted with the declarations and totaling 1139 pages provided the examiner with the facts that were allegedly misrepresented and thus, the declarant’s statements are immaterial when taken in this context.

*6 BrokerTec points to Paul’s declaration that it contends includes demonstrably false statements that are inherently material. BrokerTec also points to Fraser’s declaration as providing a broad definition of “new rules” encompassing conditional prompting. BrokerTec further contends that it does not matter whether the “new rules” portion of the Super System was within the prior art because declarations are, by their very nature, highly material. Responding to Cantor’s assertion that the statements made in the declaration were not material in light of the exhibits submitted with the declarations, BrokerTec contends that by submitting portions of the source code and misrepresenting what portions of the code included, Cantor “left the examiner with the impression” that no further investigation was necessary. *See Semiconductor Energy Lab. Co. v. Samsung Elecs. Co.*, 204 F.3d 1368, 1377 (Fed.Cir.2000).

[8] Under the reasonable examiner standard, information is material when “a reasonable examiner would consider it important in deciding whether to allow the application to issue as a patent.” *A.B. Dick Co. v. Burroughs Corp.*, 798 F.2d 1392, 1397 (Fed.Cir.1986); *see also Digital Control Inc. v. Charles Mach. Works*, 437 F.3d 1309, 1316 (Fed.Cir.2006) (suggesting that the reasonable examiner standard remains the starting point for assessing materiality).

[9] False statements are more likely material when embodied in declarations or affidavits submitted to the PTO. *See Refac Int'l, Ltd. v. Lotus Dev. Corp.*, 81 F.3d 1576, 1583 (Fed.Cir.1996); *Rohm & Haas Co. v. Crystal Chem. Co.*, 722 F.2d 1556, 1571 (Fed.Cir.1983). This court has “previously found that the submission of a false affidavit may be determined to be ‘inherently material.’ ” *Digital Control*, 437 F.3d at 1318. In *Ferring B.V. v. Barr Laboratories Inc.*, we concluded that the failure to disclose possible bias of the declarants constituted a material omission where the declarations were submitted to

--- F.3d ----

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

(Cite as: --- F.3d ----)

Page 8

overcome a prior art rejection. 437 F.3d 1181, 1190 (Fed.Cir.2006); *see also Refac*, 81 F.3d at 1581 (finding declarations submitted to establish enablement but failing to disclose possible bias of declarants were material). Similarly, in *Digital Control*, we concluded that false statements made in a declaration to antedate prior art were material. 437 F.3d at 1318.

The issue of materiality here concerns declarations that Cantor submitted to disclose the Super System and attempt to cure the earlier non-disclosure of this system. One of the declarations stated “[t]he Super System ... did not include new rules” and that “[t]he Super System code was based on ‘old rules’ in which each successive broker had a period of exclusive control over the trade.” Decl. of Bijoy Paul, ¶¶ 11, 20.

Cantor first argues that the statements were not material because there was no evidence that the new rules code was in actual use in its trading rooms. With respect to this argument, the district court found that Mr. Paul’s statement was false and that “[w]hether this code was activated or not, the applicants knew that the Super System contained the [new rules] code.” *Unenforceability Ruling*, 417 F.Supp.2d at 598. That the Super System included a form of “new rules” would surely have been important to a reasonable examiner in deciding whether to allow the 526 application to issue as a patent-in fact, Cantor acknowledges this in its brief to this court. *Brief of Appellant* at 27, *eSpeed, Inc. v. BrokerTec USA, L.L.C.*, No. 06-1385 (July 14, 2006) (“At most, Super System’s unused code for trading states may represent prefilig experimentation or development that could have been deemed material to the 733 Application under the broadest, ‘reasonable examiner’ standard.”). Thus, there is no question that the statements at issue here are material. *Accord Pharmacia*, 417 F.3d at 1373 (“[T]hese misleading declarations go to the very point of novelty.”).

*7 Cantor’s reliance on *Juicy Whip, Inc. v. Orange Bang, Inc.* is misplaced. 292 F.3d 728 (Fed.Cir.2002). In *Juicy Whip*, the patentee appealed a final judgment holding the patent in suit unenforceable based, in part, on the submission of two declarations. *Id.* at 731, 733. We reversed. *Id.* at 731. With respect to the first declaration, the record showed that

there were no false statements in the declaration and that the examiner misunderstood the declaration. *Id.* at 733, 744 (“[I]t is undisputed that every statement in the declaration is a true statement.”). Despite the examiner’s misunderstanding of the declaration, it was submitted on two other occasions to the PTO and each time the examiner’s misunderstanding was not corrected. *Id.* at 744. With respect to the second declaration, there was no proof of intent on the record and a finding of inequitable conduct could not be sustained. *Id.* at 745. The present case is markedly different than the *Juicy Whip* case. Here, following a bench trial, the district court concluded that the declarations include false statements and were worded in such a way as to deceive the examiner.

We also reject Cantor’s second argument that the examiner was on notice regarding the existence of the “new rules” in the Super System because of the submission of source code and functional and design specifications to the PTO with the declarations. This case is similar to *Semiconductor Energy Laboratory*, where we upheld a finding of inequitable conduct based on the submission of a partial translation of a Japanese prior art reference and a concise statement regarding that reference even though the entire Japanese language reference was submitted to the PTO. 204 F.3d 1368. Although parts of the reference were translated, other, more material sections of the reference remained untranslated. *Id.* at 1372. We found that “[b]y submitting the entire untranslated ... reference to the PTO along with a one-page, partial translation focusing on less material portions and a concise statement directed to these less material portions, [the applicant] left the examiner with the impression that the examiner did not need to conduct any further translation or investigation.” *Id.* at 1377. Even though Cantor provided the examiner with 1139 pages of material describing Super System and its early modifications, we agree with the district court that the “blizzard of paper” submitted to the PTO in conjunction with the declaration stating that the Super System did not include “new rules,” *Unenforceability Ruling*, 417 F.Supp.2d at 598, “left the examiner with the impression that the examiner did not need to conduct any further ... investigation,” *Semiconductor Energy Laboratory*, 204 F.3d at 1377, including an ana-

--- F.3d ---

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

(Cite as: --- F.3d ----)

Page 9

lysis of the portions of Super System source code provided by Cantor to the PTO. Given the foregoing, the district court's conclusion that the false statements in the declarations submitted to the PTO were material is not clearly erroneous.

B. Intent to Deceive

*8 [10][11] To satisfy the intent to deceive element of inequitable conduct, "the involved conduct, viewed in light of all the evidence, including evidence of good faith, must indicate sufficient culpability to require a finding of intent to deceive." *Kingsdown*, 863 F.2d at 876. There is no requirement that intent to deceive be proven by direct evidence; in fact, it is rarely proven by such evidence. Intent to deceive may be "inferred from the facts and circumstances surrounding the applicant's overall conduct." *Impax Labs. v. Aventis Pharms.*, 468 F.3d 1366, 1375 (Fed.Cir.2006) (citing *Merck & Co. v. Danbury Pharmacal, Inc.*, 873 F.2d 1418, 1422 (Fed.Cir.1989)).

[12] An inference of intent may arise where material false statements are proffered in a declaration or other sworn statement submitted to the PTO. For example, in *Paragon Podiatry Laboratory, Inc. v. KLM Laboratories, Inc.*, this court stated that the inference that material false statements contained in an affidavit were made with deceptive intent "arises not simply from the materiality of the affidavits, but from the affirmative acts of submitting them, their misleading character, and the inability of the examiner to investigate the facts." 984 F.2d 1182, 1188 (Fed.Cir.1993); see also *Refac*, 81 F.3d at 1583 ("The affirmative act of submitting an affidavit must be construed as being intended to be relied upon."); *Rohm & Haas*, 722 F.2d at 1571 (stating that the submission of affidavits "usually will support the conclusion that the affidavit in which [the material false statements] were contained was the chosen instrument of an intentional scheme to deceive the PTO").

Cantor contends that the district court improperly inferred intent from the materiality of the declarations. Cantor also contends that the district court ignored evidence that the applicants acted in good faith by submitting portions of the source code and functional

and design specifications for the software along with their declarations. Furthermore, Cantor points to the patent examiner's consideration of each of the exhibits to the declarations. BrokerTec argues that Cantor did not explain the relevance of the snippets of new rules code that can be found by digging through the 1139 pages of exhibits; worse, Cantor intentionally obscured those teachings by the deceptive, and in one case, an outright false statement in the declarations. BrokerTec also argues that the patent examiner's consideration of the exhibits to the declaration has no bearing on the question of intent.

[13][14] We conclude that the district court's finding of intent is not clearly erroneous. The applicants submitted the declarations at issue in an apparent attempt to purge possible inequitable conduct in the 733 application and to disclose the Super System in the 526 application. Instead of being candid, Paul's declaration disingenuously states that the Super System did not include new rules. The district court was free to draw an inference that these declarations were "the chosen instrument of an intentional scheme to deceive the PTO," *Rohm & Haas*, 722 F.2d at 1571, because "[t]he affirmative act of submitting an affidavit must be construed as being intended to be relied upon," *Refac*, 81 F.3d at 1583.

*9 We disagree with Cantor's contention that the district court failed to consider evidence of good faith. The district court acknowledged Cantor's argument, but disregarded Cantor's view of the evidence. *Unenforceability Ruling*, 417 F.Supp.2d at 598. We see no clear error in its conclusion.

III. CONCLUSION

Based on threshold findings of materiality and intent to deceive, the district court balanced the equities and found that the 580 patent is unenforceable. The district court did not abuse its discretion in so holding. Because the 580 patent is unenforceable, we need not reach the other issues raised by Cantor on appeal. Therefore, the judgment of the district court is

AFFIRMED.

No costs

--- F.3d ----

Page 10

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

(Cite as: --- F.3d ----)

FN1. These other trading states, which are of less relevance to our decision, include a “Bid/Offer” state and an “Uncleared Bid/Offer” state.

FN2. During oral argument, Cantor's counsel admitted that the code for new rules—including the new workup rules—was accessible to brokers using the Super System and that the code did not need to be modified to enable the use of the new rules.

FN3. The declaration of Stuart A. Fraser indicated that “[t]he Super System did not use the ‘new rules.’ ” Decl. of Stuart A. Fraser, 526 application, ¶ 21 (Jan. 31, 2002).

C.A.Fed. (Del.), 2007.

eSpeed, Inc. v. BrokerTec USA, L.L.C.

--- F.3d ----, 2007 WL 817665 (C.A.Fed. (Del.))

Briefs and Other Related Documents (Back to top)

- 2006 WL 3294766 (Appellate Brief) Reply Brief of Plaintiffs-Appellants Espeed, Inc., Cantor Fitzgerald, L.P., CFPH, L.L.C. and Espeed Government Securities, Inc. (Oct. 20, 2006)
- 2006 WL 2984263 (Appellate Brief) Brief of Defendants-Appellees BrokerTec USA, L.L.C., Garban, L.L.C., OM Technology US, Inc. and OM Technology AB (Sep. 22, 2006)
- 2006 WL 2252252 (Appellate Brief) Nonconfidential Brief of Plaintiffs-Appellants Espeed, Inc., Cantor Fitzgerald, L.P., CFPH, L.L.C. and Espeed Government Securities, Inc. (Jul. 14, 2006)

END OF DOCUMENT